arXiv:1610.03724v1 [stat.AP] 12 Oct 2016

The Weak Efficient Market Hypothesis in Light of Statistical Learning

L. Fiévet and D. Sornette

October 13, 2016

Chair of Entrepreneurial Risks, ETH Zürich, Scheuchzerstrasse 7 (SEC F), CH-8092 Zürich, Switzerland

Abstract

We make an unprecedented evaluation of statistical learning methods to forecast daily returns. Using a randomization test to adjust for data snooping, several models are found statistically significant on the tested equity indices: CSI 300, FTSE, and S&P 500. A best Sharpe ratio portfolio has abnormal returns on the S&P 500, breaking even with the market at 10 bps in round trip costs. The returns produce statistically significant intercept for factor regression models, qualifying as a new anomalous 3-day crisis persistency factor. These results open the path towards a standardized usage of statistical learning methods in finance.

Extensive research effort has been dedicated to forecasting stock market returns using conventional regressive models (Atsalakis and Valavanis, 2013), unconventional methods (Atsalakis and Valavanis, 2009), and technical trading rules (Hsu et al., 2010). Conventional regression methods typically assume a rigid functional form, which limits the return dependencies they can capture. In contrast, unconventional methods can have an arbitrary level of flexibility to capture dependencies, however their calibration requires non-standard statistical theory (Mills, 1999, p. 224). This issue has led the research on unconventional methods to use a variety of performance metrics that are difficult to compare across studies. Universal metrics, such as a test statistic with respect to a null distribution, or the break-even transaction cost, are rarely provided. This lack of standardized performance metrics for unconventional methods has held back their usage in econometrics.

While the usage of unconventional methods remains limited, the wide use of technical trading rules, many of which are not calibrated using standard statistical theory, has led to the development of tests to determine their statistical significance. The Reality Check method by White (2000) achieved this feat, and has led to several extensions by Romano and Wolf (2005), Hansen (2005), and Hsu et al. (2010). These frameworks provide a robust methodology to compute the statistical significance of a model, adjusting for data snooping in a universe of models. Unfortunately, these frameworks remain insufficiently leveraged to determine the statistical significance of interesting results, such as the findings of Niederhoffer and Osborne (1966), Zhang (1999), Leung et al. (2000), Andersen and Sornette (2005), Zunino et al. (2009), James et al. (2014) and Satinover and Sornette (2012a,b), which show departure from the martingale condition for daily returns on a variety of assets.

The Reality Check method computes the statistical significance of a model, but is insufficient to confirm or reject the weak efficient market hypothesis (denoted by EMH in this paper; Fama, 1970, 1991), as it is unclear how to select that model ex-ante. Timmermann and Granger (2004) provide an operational definition of the EMH, requiring a search technology that would have selected ex-ante a profitable model from the universe of models. In case the EMH is rejected, the anomalous returns have to be tested with the three-factor and four-factor regression models, by Fama and French (1993), respectively Carhart (1997). These models incorporate all the known anomalies to the EMH, and determine if the observed anomalous returns can be explained by already known factors.

This paper implements an unprecedented evaluation of statistical learning methods (Hastie et al., 2001) to forecast daily returns of equity indices, using a randomization test based on the framework of Romano and Wolf (2005, 2016). Each method gives rise to several models, which are parametrized by the number of lags and in-sample length used for calibration. These methods are particularly attractive to find non-linear return dependencies, which could go unnoticed in parametric regression models. The models are evaluated by computing the p-values adjusted for data snooping of three performance metrics, namely the directional accuracy based on the non-parametric tests by Pesaran and Timmermann (1992), the compounded wealth of the entire trading period, and the Sharpe ratio. To evaluate the trading

performance achievable ex-ante, we use a best Sharpe ratio search technology. Using factor regression models, the obtained returns are tested for correlation with known anomalous factors. Assuming transaction costs constant in time, and accounting for the monthly risk free rate, the break-even cost with the buy and hold strategy is computed.

The research by Fernández-Delgado et al. (2014) showed that a small number of statistical learning methods provides the highest performance across a large range of datasets. This motivated the construction of our universe of models from these top performing methods. For each method, the parameters used to generate different models is carefully discussed. Additionally, the GARCH model is used as an econometric reference model. This led to a total of 3136 uniquely defined models. Given this large universe of models, a mean variance portfolio as search technology revealed technically challenging to implement (Senneret et al., 2016). Therefore, we chose the best Sharpe ratio search technology, which invests at every time step into the model with the best past Sharpe ratio.

Statistical learning methods describe many decision boundaries that cannot be captured with linear models. For example, the logical exclusive-OR function (XOR) is not linearly separable, but can be well approximated with a decision tree. This is especially promising because the heuristics used by humans (Tversky and Kahneman, 1974) are likely to include linearly inseparable functions, patterns which are much better captured by decision trees. Unfortunately, most of the statistical learning methods cannot be calibrated using standard statistical tools, and require the use of model free performance metrics. Following the studies of White (2000) and Sullivan et al. (1999), we chose to benchmark directional accuracy and the Sharpe ratio as a proxy for risk adjusted returns. The directional accuracy is particularly relevant for our study, as many models work on binary returns of up and down moves, and may show abnormal directional accuracy but normal Sharpe ratio. While the Sharpe ratio could be used as a measure of profits, needed to test the EMH, it remains a proxy that could be biased, and we therefore use compounded wealth to benchmark profits.

We conduct our empirical study on daily returns of the CSI 300 during the period 2005-2015, and the FTSE and S&P 500 during the period 1995-2015. The following confidence levels were obtained when adjusting for data snooping across the entire model universe. The directional forecast accuracy is rejecting the null hypothesis at the 99.8% level for the CSI 300 and the S&P 500, and at the 92.8% level for the FTSE. The compounded wealth is rejecting the null hypothesis at the 91% level for the CSI 300, and at the 97.5% level for the FTSE and S&P 500. The p-values obtained for the Sharpe ratio were mostly identical to the compounded wealth p-values. The best performing models break even with the buy and hold strategy at transaction costs per round trip of 28.5 bps for the CSI 300, 15.9 bps for the FTSE, and 12.9 bps for the S&P 500. These profitability thresholds clearly exceed the 5 bps transaction costs available nowadays on the future market.

The best Sharpe ratio search technology, used to determine the ex-ante performance, would not have been profitable on the CSI 300, but breaks even with the buy & hold strategy at 2.2 bps per round trip for the FTSE, and 9.9 bps for the S&P 500. We regress

the search technology returns for the S&P 500 with the three- and four-factor models at zero transaction costs. The monthly intercept at 0.49% is significant with a test statistics of $t = 2.76$, and remains significant at $t \geq 2$ up to round trip costs of 3.6 bps. The observed returns are therefore unexplained by current factors, and constitute a new anomalous 3-day crisis persistency factor (or stylized fact) for the S&P 500. Our time trend analysis uncovered that the most profitable periods happened during the volatile crashes of the dot-com bubble, the financial crisis, and the European debt crisis.

This paper continues as follows, Section I constructs methodically the universe of statistical learning models that can meaningfully be calibrated, and discusses the rational for the chosen search technology. Section II presents the performance metrics used to benchmark the models, as well as the randomization test used to computed p-values adjusted for data snooping. Section III presents the empirical results, while the last section concludes and outlines a path for future research.

I. Statistical Learning and the EMH

A. Definitions

Predicting new observations based on past observations is the common denominator of empirical sciences. During past centuries, each science has developed its own terminology and models to describe its observations. For a number of simple systems, such as predicting the periodicity of a pendulum, an exact solution could be derived from first principles. However, many complex systems, such as gene expression, wages, or stock markets discussed by James et al. (2014), escape the realm of full predictability due to the large amounts of data and variables involved. Nonetheless, with the advent of cheap computation, a large number of generic methods, referred to as **statistical learning** (or **machine learning** in computer sciences), have been developed to extract predictability from such systems. These methods are discussed in detail by Hastie et al. (2001) and James et al. (2014).

At an abstract level, the problem is to estimate the systematic relation f between a set of **inputs** X and dependent **outputs** Y, such that

$$Y = f(X) + \epsilon. \tag{1}$$

The error term ϵ has mean zero, and is independent of X. To illustrate, the output could be the wage of a person, and the inputs are some characteristics of that person, such as age and years of education. Determining the function f from first principles is difficult or sometimes impossible, and could introduce the modeler's bias. To overcome these issues, a statistical learning approach learns the function f based on a **training dataset**

$$\mathcal{D} = \{(X_1, Y_1), \ldots, (X_N, Y_N)\} \tag{2}$$

with a sufficient number of samples. In the wage example, a **sample** would be a data tuple $(X = (age, education), Y = (wage))$ associated to a person. The set of statistical learning methods include parametric and non-parametric methods, ranging from low flexibility (e.g. linear regression) to high flexibility (e.g. support vector machine with non-linear kernels). The challenge in the estimation of f for flexible methods is to avoid over-fitting the error term associated with the dataset $\mathcal{D}$. The method has to only extract the systematic information, so that the estimated function works on any dataset $\mathcal{D}'$ with new samples of the same type.

To use statistical learning within the context of the efficient market hypothesis (EMH), a precise definition of the latter is needed. The EMH has been defined in a robust way as follows.

> **Efficient Market Hypothesis:** "A market is efficient with respect to the information set Ω_t, search technologies S_t, and forecasting models M_t, if it is impossible to make economic profits by trading on the basis of signals produced from forecasting model in M_t defined over predictor variables in the information set Ω_t and selected using a search technology in S_t." Timmermann and Granger (2004)

This definition is fully operational, nonetheless we consolidate two aspects. First, it is unclear how to determine the profitable search technology in the set S_t. This would require a single top level search technology S_t^*, selecting a search technology in S_t. To avoid such ambiguity, we limit this study to a single search technology. Second, the economic profits must be significant at the 95% level (or higher) in comparison to the profit distribution of random models, after adjusting for multiple testing across several assets. Otherwise, any spurious economic profits found in a single asset would suffice to claim inefficiency.

We remark that the above definition implicitly assumes that the set of forecasting models includes all evaluated models. An ex-post selection of models could falsify the results by introducing a bias towards the top performing models. Likewise, the search technology should be determined before evaluating the models to avoid any ex-post biases towards some models.

Many financial forecasting models explicitly incorporates knowledge (information) about the market structure underlying an asset price. However, a strict interpretation of the weak(est) EMH forbids any information besides the past asset returns, hence excluding knowledge of the market structure. Therefore, the only option to forecast asset returns is to test a set of statistical learning methods free of implicit financial knowledge. This includes linear models and extensions thereof. However, the top performing methods can provide information about the economic structure generating the asset returns.

B. Parametrized Map of the Information Set

We now consider the weak(est) EMH, where the information at time t is given by all past returns $\Omega_t = \{r_1, \ldots, r_t\}$. The information set Ω_t in its raw form is unsuitable to calibrate

a statistical learning method. To create a set of samples with inputs and outputs, the past returns Ω_t are sliced into $t-1$ sets of past returns and their following returns $\bar{\Omega}_{\tau\in[1,\,t-1]} = \Omega_t \backslash \Omega_\tau$ as

$$\mathcal{D}_t = \left\{ \left(\Omega_1,\, \bar{\Omega}_1\right),\, \ldots,\, \left(\Omega_{t-1},\, \bar{\Omega}_{t-1}\right) \right\}. \tag{3}$$

The defined dataset $\mathcal{D}_t$, constructed from Ω_t, matches exactly the structural requirement defined in Equation 2 for a statistical learning approach. Every time step becomes a sample, with all the preceding returns as inputs and all the following returns as outputs.

Although the sample size should be maximized to avoid over-fitting, it may not be optimal to keep all past returns as the distribution of returns is usually not strictly stationary (e.g. see Mills (1999)). For example, the recursive regression of an ARCH type model was used by Phillips et al. (2011) to determine structural breaks in the NASDAQ, where a regime change occurs. The regression of an ARCH model on different regimes yields fundamentally different regression factors. Such non-stationarity could decrease the predictive power of past returns over time or even negatively impact performance after a regime change. To minimize the mixing of different regimes, we limit the number of samples used for calibration to a constant in-sample length l_{is}. As we are interested in short term daily dependencies while allowing for changing market regimes, we bound the in-sample length to two years ($l_{is} \leq 500$ trading days). To avoid almost identical models, and reduce computation time, we explore the in-sample length in steps of 10 days. As we are interested in forecasting daily return one day ahead, we limit the output space to the first return following the inputs.

In Equation 3, the inputs of a sample have been defined as all past returns available at time t, causing the input space to increase over time. An input space increasing in size suffers from the curse of dimensionality and must be reduced. Assuming binary inputs and outputs, and considering only m lags as inputs at every time step, there are 2^m distinct input sequences. For each input sequence, there are two different possible outputs, whose probabilities are determined based on the training data. To distinguish which one of the two outputs is more likely, with a granularity of 5%, at least 20 samples are required. This leads to the inequality condition $20 \times 2^m \leq l_{is}$, which ensures that on average twenty instances of each sequence of length m are present in the time window of length l_{is}. Assuming at most two years of calibration data ($l_{is} \leq 500$), this condition implies $m \leq 5$. As the input space of the regression problem is more fine grained, this same bound on the lag applies to the regression methods as well. As a single lag is trivial, we explore the lags in the set $m \in \{2,\, 3\,,\, 4,\, 5\}$. The smallest lag $m = 2$ implies the lower bound $l_{is} \geq 80$, however, for completeness (or curiosity), we start at an in-sample length of $l_{is} = 20$, amounting to a total of 49 possible lengths $l_{is} \in [20,\, 30,\, \ldots,\, 500]$.

Given the restrictions discussed, the sample set at time t is now written as

$$\mathcal{D}_t = \left\{ \left(\Omega_{t-l_{is},\,m},\, r_{t-l_{is}+1}\right),\, \ldots,\, \left(\Omega_{t-1,\,m},\, r_t\right) \right\} \quad \text{where} \quad \Omega_{t,\,m} = \left\{ r_{t-m+1},\, \ldots,\, r_t \right\}. \tag{4}$$

It should be noted that many other choices can be made for the inputs and outputs. Based

on all inputs Ω_τ and outputs $\bar{\Omega}_\tau$ one can for example compute trends on several days, or volatility levels, and predict them using the same approach.

C. Regression & Classification

Outputs can be grouped into two types, namely **quantitative** outputs such as wages, and **categorical** outputs such as gender. Statistical learning methods predicting quantitative outputs are referred to as **regressor**, and those predicting categorical outputs as **classifiers**.

Asset returns are quantitative outputs that can be predicted using regressors. The degrees of freedom in regressors have a variety of origins: the parameters of the function to regress, for example the coefficients of a linear regression; the parameters of a kernel transforms of the inputs, for example the exponent in a homogeneous polynomial kernel; and the number of partitions of the input space, for example the leaves in a regression tree. With too few degrees of freedom, the estimated regressor will still suffer from a high bias, failing to model relevant relations between inputs and outputs. With too many degrees of freedom, an estimated regressor will have high variance, modeling the random noise in the training data. In the case of asset returns, the random noise component is large and therefore the optimal methods should have a low number of degrees of freedom to avoid high variance. Typically, the optimal tradeoff between bias and variance can be found using a test set and cross-validation techniques. However, these techniques are not applicable to returns due to there time dependence (see Appendix A for a discussion), leaving us without a generic approach to constrain the number of degrees of freedom.

Given this unavoidable bias-variance dilemma, we turn towards preprocessing the data in ways that reduce variance while retaining the meaningful information. In particular, to create a successful trading strategy, a forecast does not need to predict the exact return. Merely predicting the daily trend is sufficient to generate profits. Likewise, from a behavioral perspective of the agents acting on stock markets, one could argue that the agents do not care about the precise value of returns. Fundamentalists trade based on their fundamental analysis of companies, and chartists mostly look at trends and a finite set of patterns. Consequently, it is meaningful to map returns onto a set of categorical outputs that relate closely to the discrete mental models used by the agents. Tversky and Kahneman (1974) showed extensively that humans, even when trained in statistics, often rely on heuristics to answer statistical questions. Further on, Gary and Wood (2010) studied human mental models in solving strategic problems, and reported that good heuristics often suffice to achieve good performance. Within this heuristic perspective, the agents trading on stock markets make the strongest distinction between negative and positive returns. Therefore, the most sensible preprocessing is a binary map of the outputs to down and up moves. Such binary outputs can be predicted using classifiers. In a second step, the same binary preprocessing can as well be applied to the inputs.

Beyond binary mappings, there are a large number of trinary and quaternary mappings

that are meaningful and can be calibrated robustly with the available number of samples. As well, there is no mathematical constraint requiring to apply the same mapping to the inputs and outputs. The mapping applied to the inputs can even be defined so as to map differently the individual components of an input. Nevertheless, the more complex mappings are partially described by the binary mapping and regression methods applied to the quantitative returns. Therefore, and in order to avoid computation overload, we will study here only models using the binary mapping.

D. Constructing the Universe of Models

At the time of writing, there are hundreds of statistical learning methods available in the literature (Fernández-Delgado et al., 2014). In James et al. (2014), several of the most common methods have been tested on returns of the S&P 500 between the year 2001 and 2005, with the highest performing method being the quadratic discriminant analysis. The present paper evaluates the set of methods presented in Table I, covering the families of linear regression (GARCH), discriminant analysis, logistic regression, nearest neighbors, support vector machines, decision trees, and bagging and boosting of decision trees. These are the most commonly used methods and known to be among the top performers across a large variety of problems. Neural networks have been excluded, as they are computationally much more intensive and difficult to interpret.

Classification methods (and most regression methods) can be visualized as separating the samples with decision boundaries maximizing the predictability within each group formed by the boundaries. The difference between the methods mostly lies in the shape of the decision boundaries. For the reader unfamiliar with statistical learning, some visual examples are provided in Figure 1. These figures represent a two dimensional input and a three class categorical output (e.g. negative returns, small returns and positive returns). Subfigure a) shows the optimal decision boundary from a linear discriminant analysis separating three multivariate normal distributions. Subfigure b) shows some optimal hyperplane boundaries used in support vector machines. Finally, subfigure c) represent the boundaries of a decision tree obtained by recursively splitting the dataset.

The selected regression methods are applied to the quantitative inputs and outputs. The classification methods are applied to quantitative inputs but binary outputs of down and up moves. The case of binary inputs organizes the samples neatly into m dimensional hypercubes as illustrated in subfigure d). In such a configuration, the shape of the decision boundaries becomes irrelevant and all methods will make the same prediction. An input space partitioned into hypercubes can be exactly matched by the decision boundary of decision tree based methods. Therefore, binary inputs are predicted using a decision tree. The prediction of a decision tree is the majority vote (classification) or mean (regression) within each hypercube. Table II provides an overview of all the 16 possible combinations of using regressors and classifiers on quantitative or binary inputs and outputs. Given the four differ-

Method	Description	Parameters
Regressors		
GARCH(p, q)	Linear regression model, with a nested linear regression model for the variance of the error terms.	Lags $p = q = m$
Classifiers		
Linear **D**iscriminant	Assumes **k** classes from independent multivariate distributions. Computes the optimal linear boundary between the distributions.	No shrinkage
Quadratic **D**iscriminant	Assumes **k** classes from independent multivariate distributions. Computes the optimal quadratic boundary between the distributions.	
Logistic **R**egression	Computes the logistic decision boundary between two classes.	
Regressors/Classifiers		
Nearest **N**eighbors	Finds the **k** nearest neighbors of a data point and computes their mean output or majority vote.	$k = 5$, Minkowski metric, uniformly weighted
Support **V**ector **M**achine	Computes the maximum-margin separating hyperplane.	Linear kernel
Decision **T**ree	Partitions the features space by recursive splitting at optimal values along a single feature. Computes the mean or majority vote in each subset.	Any depth, minimum one sample per leaf
Random **F**orest	Averages the prediction of many decision trees trained on subsamples of the training data to avoid over-fitting.	10 trees
Gradient **B**oosting	Recursively creates decision trees to reduce the remaining errors.	Max depth 3, 100 boosts, learning rate 0.1

Table I: **Description and relevant parameters of the statistical learning methods evaluated in the present work to forecast daily returns.** The table separates regression only, classification only (e.g. down/up moves), and dual purpose methods.

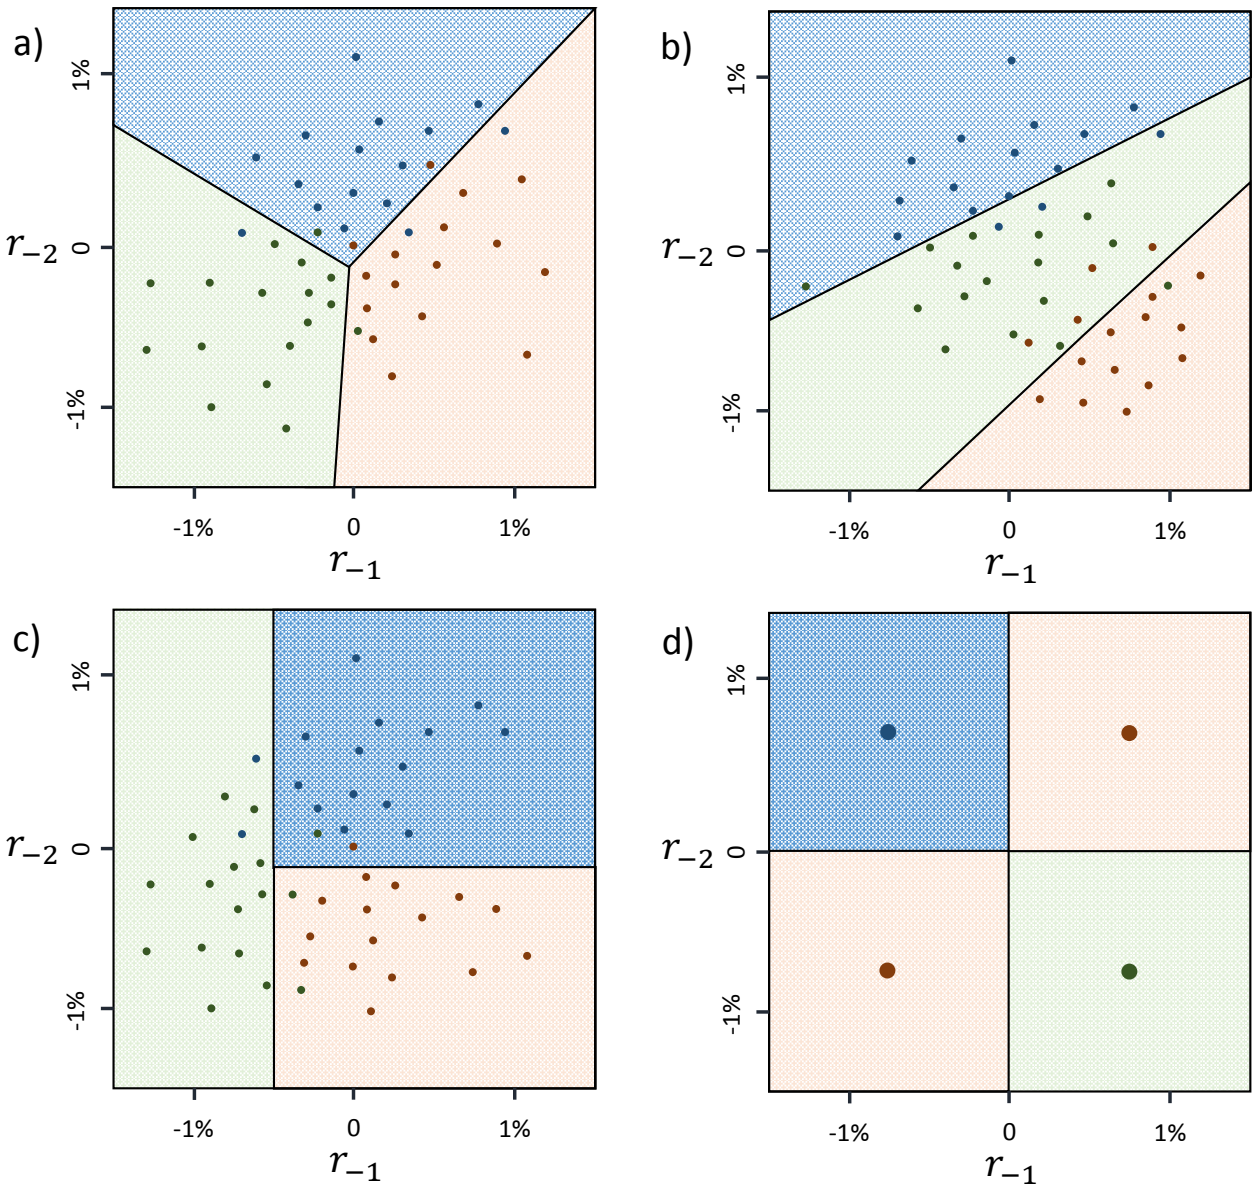

Figure 1: **Schematic examples of decision boundaries for several statistical learning methods fitted on the inputs** (r_{-2}, r_{-1}) **of two past quantitative returns and three possible output classes.** The methods are: a) linear discriminant analysis; b) support vector machine; c) decision tree; d) decision tree fitted on binary returns (the choice of having twice the red class as the majority vote is arbitrary).

Inputs	Outputs	Methods
Quantitative	Quantitative	6 Regressors
Quantitative	Binary	8 Classifiers
Binary	Quantitative	1 Regression Tree
Binary	Binary	1 Decision Tree
	Total:	16

Table II: **Summary of all the tested combinations of quantitative and binary inputs and outputs.** The number of available classifiers and regressors is taken from the selection presented in Table I.

ent lags $m \in \{2, 3, 4, 5\}$ and the in-sample length $l_{is} \in [20, 500]$ in steps of 10, our universe M_t has $4 \times 49 \times 16 = 3136$ uniquely defined models. In this study, the set of models is time independent, and we write simply M.

E. Models Discussion & Trading Signal

To make profits (or losses), a model's predictions have to be converted into a signal s_t determining the market position to take. The simplest conversion is to take a long position ($s_t = 1$) upon an up move prediction and a short position ($s_t = -1$) upon a down move prediction. In case short positions are not allowed, a down move prediction would require to stay out of the market ($s_t = 0$). However, many models do not simply predict an up or down move but provide a probability or amplitude for the predicted return, and it would be unreasonable to treat all predictions equally. One could take a position proportional to the excess probability, or stay out of the market when the probability is too uncertain. To reduce a strategy's turnover, a change in position could be allowed only after two consecutive predictions in the opposite direction. In short, the possibilities to create trading signals are endless, and testing several trading signals would multiply the size of our model universe. Such an additional exploration of trading signals is not at the heart of our study, which aims to compare the performance of different statistical learning methods. Hence, we will refrain from testing multiple trading signals and aim at defining a single trading signal that allows for a fair comparison between the models.

The GARCH model captures linear dependencies between lagged returns, and linear dependencies between the error terms and volatility. At each step, the forecasted return and volatility level are used to compute the probability of an up or down move. The linear

regression of highly stochastic returns tends to simply capture the current market trend, and the resulting strategies are highly correlated with trend following.

The linear and quadratic discriminant analysis can determine if the m-variate distribution of returns followed by a down move significantly differs from the m-variate distribution of returns followed by an up move. Whenever the two distributions are distinguishable, the prediction probability will rise above randomness. The largest difference between the two distributions is located in the tails, where the predictions with the highest probabilities will be made. The tails of the distributions are associated to large returns, consequently the discriminant methods could detect predictability linked to high volatility.

The logistic regression of returns can detect if any direction in the input space, with respect to a reference point, is dominantly followed by a down or up move. Due to the approximate normal distribution of returns around zero, the reference point will be found near zero. Predictions with high probability will be made following large returns, consequently the logistic regression can detect predictability linked to high volatility in a certain direction in the input space.

The nearest neighbor method is evaluated for $k = 5$, computing for a given input the mean or majority vote of the five nearest samples (using an euclidean metric) in the training set. In the classification case, the odd number of neighbors ensures a clearly defined outcome. The choice of only five neighbors guarantees the estimation of a neighborhood close to the evaluated input even in the case of $m = 5$ lags. This method can detect the overall market response predictability to a given sequence of past returns. Predictability in specific neighborhoods can go unnoticed if diluted by other neighborhoods where the market response is reverting regularly.

The support vector machine (SVM) method separates the features space by a hyperplane, maximizing the class asymmetry between the two areas defined by the hyperplane. This method is ideal to detect any imbalance of the distribution of up or down moves in the input space. On the downside, the standard SVM classifier does not support the computation of a prediction probability.

The decision tree method partitions the input space into disjoint n-orthotopes (hyper-rectangles), maximizing the proportion of a single class in each of them. Allowing the tree to be refined to an arbitrary depth is ideal to capture small pockets of predictability. While we want to test for such pockets, this leads to an almost certain over-fitting that is addressed by bagging and boosting methods. The bagging method, called a random forest, averages multiple trees calibrated on subsets of the training data to reduce the variance. The boosting method recursively reduces the error by adding new trees calibrated on the remaining error. To avoid a perfect fit from the beginning, the tree depth is limited to the number of lags in the boosting case.

In all regression cases, the prediction probability is computed by evaluating the cumulative distribution function of the past returns at the predicted value. For all classifiers, the probability is defined by the method or otherwise set to one. To avoid penalizing methods

with many uncertain predictions, we introduce a threshold ρ_τ below which a trading strategy stays out of the market. Adapting the threshold dynamically, based on the method and the number of samples, is unfortunately out of reach given the limited control of the Scikit-learn package (see Pedregosa et al. (2011) for an overview of the package). Therefore, we decided to use a lower bound based on the binary decision tree predictor with the fewest lags ($m = 2$) and the median in-sample length $l_{is} = 250$. This method has four possible input states (00, 01, 10, 11), with on average $250/4 \sim 62$ samples per input, implying a minimum granular probability larger than 50% equal to $32/62 \sim 51.6\%$. Consequently, we set $\rho_\tau = 0.515$ as a threshold for all methods to eliminate biases by too uncertain predictions.

In the case of binary classification, a systematic bias is introduced when the time series contains more down or up moves. In many statistical learning applications, when the imbalance is strong, a re-balancing has to be performed to produce meaningful results. However, in the case of asset returns, the imbalance is weak and can depend on the market regime (e.g. more up moves during a bull market). Therefore, we decided not to re-balance so the methods can follow trends in a similar manner as a linear regression model.

F. A Practicable Search Technology

The last piece needed to evaluate market efficiency is the search technology. Let us denote by $\vec{r}^M(t)$ the set of returns at time t of the models in the universe M of models. Then a search technology is a function $\vec{S}(M, t)$ computing the weights of a portfolio of models in M at time t, resulting in portfolio returns $r^S(t) = \vec{S}(M, t-1) \cdot \vec{r}^M(t)$. A mean variance portfolio of models, based on their past returns, would provide an excellent solution of the search process. However, to invert with reasonable errors the covariance matrix of the returns of all models, the number of past returns must be greater than or equal to the number of models in M (see Senneret et al. (2016) for a discussion). Given the 3136 models, this would require at least 12 years of past returns for all models, a time span too long for the present work. Instead of the mean variance portfolio, we select at every time step the single model with best in-sample risk adjusted returns, using the Sharpe ratio (Sharpe, 1994).

To determine an appropriate window size for computing the best Sharpe ratio, we consider the worst case scenario where, out of the 3136 models, all are random except one. We base our computation on a sequence of N returns drawn in $\mathcal{N}(0, \sigma_r = 0.01)$, zero risk free rate, and random models predicting up or down with a probability of 50%. The choice of $\sigma_r = 1\%$ corresponds to the typical daily volatility of equity indices. As the random models are always long or short, it follows that their returns are drawn in the distribution $\mathcal{N}(0, \sigma_r = 0.01)$. Consequently, the mean returns of the random models after N steps are distributed as $\mathcal{N}\left(0, \sigma_r/\sqrt{N}\right)$, and their daily Sharpe ratios are distributed as $\mathcal{N}\left(0, \sigma_{SR_d} = \sqrt{1/N}\right)$. The daily Sharpe ratios of the 3135 random models are all below the ensemble 95% confidence level of $3.05 \cdot \sigma_{SR_d}$, defined by the fact that the probability of all 3135 Sharpe ratios being

below $3.05 \cdot \sigma_{SR_d}$ is equal to 95%: $\Phi\left(3.05\right)^{3135} = 0.95$, where Φ is the cumulative distribution function of the standard normal distribution. A truly predictive model, with a compounded annual return $r_y = 20\%$ (or equivalently $r_d \approx 0.073\%$ daily), has a daily Sharpe ratio of $SR_d^{20\%} = \frac{r_d}{\sigma_r} \approx 0.073$. It follows that the truly predictive model will outperform at the 95% confidence level when $SR_d^{20\%} \geq \frac{3.05}{\sqrt{N}}$, which requires $N \gtrsim 1746$ trading days, or ≈ 7 years.

In order to exploit the full 7 years of out-of-sample data, this study will be relying on an expanding window with an initial length of 1 year. Denoting the Sharpe ratio of model $M_i \in M$ at time t by $SR_i\left(t\right)$, the search technology reads

$$S_i\left(M,\, t\right) = \begin{cases} 0 & \text{if } i \neq \underset{M_i \in M}{\operatorname{argmax}}\, SR_i\left(t\right) \\ 1 & \text{if } i = \underset{M_i \in M}{\operatorname{argmax}}\, SR_i\left(t\right) \end{cases}. \tag{5}$$

This search technology can be refined with the Reality Check method (White, 2000), as done by Hsu et al. (2016). The refined method computes an equally weighted portfolio of models with significant Sharpe ratios after adjusting for multiple testing. However, the present work is not about optimizing the search technology, and the best Sharpe ratio search technology is sufficient to estimate the ex-ante trading performance.

II. Performance Metrics & Statistical Significance

A. Methodology

In the lineage of the studies by White (2000) and Sullivan et al. (1999), we benchmark the directional accuracy and the Sharpe ratio. However, within the context of measuring profits, abnormal directional accuracy does not necessarily imply abnormal returns. There are models that can have below average directional accuracy, but still generate significant abnormal returns, making this metric incomplete (and sometimes misleading) as a proxy for profits. The Sharpe ratio is suitable to determine models with abnormal risk adjusted returns, nonetheless could heavily penalize models with large upside volatility but little downside risk. To avoid such shortcomings, we compute as well the compounded wealth, a truthful measure of profits devoid of aggregation biases most proxies suffer.

As mentioned in Section I, rejecting the EMH cannot simply be based on just the existence of some profits, but requires to prove that the profits are statistically significant. The p-values for the chosen performance metrics are computed using a randomization test of the trading signals, based on the framework of Romano and Wolf (2005). Our null hypothesis is that the model strategies are drawn from the set of random strategies with the same trading pattern and dependency structure. The distribution of a performance metric under this null hypothesis is computed by a Monte-Carlo simulation. The advantage of randomizing the trading signals, over Bootstrapping the model returns, is that we do not need to assume

stationarity of the underlying asset. This provides us with almost exact p-values. The efficient resampling based stepdown algorithm used to adjusted the p-values for multiple testing is defined by Romano and Wolf (2016). This algorithm maximizes the number of rejected null hypothesis without violating the familywise error rate.

To better understand the relationship between directional accuracy and compounded wealth after transaction costs, we compute an approximate relation between the two. First, this relation allows us to determine if a model's compounded wealth can be explained by its directional accuracy. Models that over-perform the expected wealth predict large returns better then their average directional accuracy, a beneficial feature. Second, we can use this formula to estimate the break-even trading cost for a model's compounded wealth.

B. Directional Accuracy

The directional accuracy measure was chosen based on the discussion of simple non-parametric tests of predictive performance by Pesaran and Timmermann (1992). These tests are based on mapping the real returns onto one of n_r categories and the predicted returns onto one of n_c categories. Then each sample in the dataset is an observation of a real and a predicted return, which can be allocated to the corresponding cell of the contingency matrix O of n_r rows and n_c columns. The real and predicted returns are independent under the null hypothesis of random strategies. For sufficiently large samples, Pearson's chi-square test can be used to measure the independence of observations (Pearson, 1900, Plackett, 1983). For contingency matrix O, with observation numbers $O_{i,j}$, the chi-squared test is defined as

$$\hat{\chi}^2_d = \sum_{i=1}^{n_r} \sum_{j=1}^{n_c} \frac{(O_{i,j} - E_{i,j})^2}{E_{i,j}}, \tag{6}$$

where

$$E_{i,j} = N\rho_{i\cdot}\rho_{\cdot j}, \quad N = \sum_{i,j}^{n_r,n_c} O_{i,j}, \quad \rho_{i\cdot} = \frac{1}{N}\sum_{j=1}^{n_c} O_{i,j}, \quad \text{and} \quad \rho_{\cdot j} = \frac{1}{N}\sum_{i=1}^{N} O_{i,j}. \tag{7}$$

The number of degrees of freedom of the contingency table is given by $n_f = (n_r - 1)(n_c - 1)$. The value $\hat{\chi}^2_d$ is used as the performance measure for directional accuracy. The independence test was selected because it is more conservative then the predictive failure test. The later only measures excess or deficit in the diagonal terms of the contingency matrix, terms that correspond to the number of correct predictions. However, in the case $n_r = n_c = 2$ of interest in this study, the two hypotheses converge asymptotically for $N \to +\infty$, where N is defined in Equation 7.

Under the null hypothesis of random strategies, disregarding trading patterns, we note

that the p-value can be computed analytically as

$$P\left(\chi_d^2 \geq \hat{\chi}_d^2,\, n_f\right) = \Gamma\left(\frac{n_f}{2}\right)^{-1} \Gamma\left(\frac{n_f}{2},\, \frac{\hat{\chi}_d^2}{2}\right). \tag{8}$$

The p-value is the probability of a directional chi-squared value χ_d^2 larger then the observed value $\hat{\chi}_d^2$. This result is useful to approximate efficiently the directional accuracy p-value, when a Monte-Carlo approach is not feasible because of computational limits.

C. Compounded Wealth & Sharpe Ratio

Compounded wealth is normalized with unit initial wealth ($W_1 = 1$). At each time step t the wealth $W(t)$ is invested based on the model's signal $s_t \in [-1,\, 1]$, compounding with the asset return r_t. When the prediction is inconclusive, below the probability threshold ρ_τ, the trader stays out of the market, moving his wealth to cash at the risk free rate r_t^f. The transaction cost $\Delta\varsigma$ is in percentage points of the change in position. For ease of computation, the transaction cost is assumed to be constant in time. For example, the round trip cost of taking a long position and then selling it is $2\Delta\varsigma$ percentage points. The wealth at time t can be expressed as

$$W(t) = \prod_{i=1}^{t-1} \underbrace{\left(1 - |s_i - s_{i-1}|\, \Delta\varsigma\right)}_{\text{transaction cost}} \underbrace{\left(1 + s_i \cdot r_i + (1 - |s_i|) \cdot r_{\mathrm{i}}^f\right)}_{\text{model return}}, \tag{9}$$

where $s_0 = 0$ is an out of the market signal before the first predicted signal s_1. The Sharpe ratio of a model is computed from the model returns $r^M(t) = \frac{W_t}{W_{t-1}} - 1$ as

$$SR(t) = \frac{E\left(r_t^M\right) - E\left(r_t^f\right)}{\sqrt{E\left(\left(r_t^M\right)^2\right) - E\left(r_t^M\right)^2}}. \tag{10}$$

The expected values are estimated from the samples. Following the results of Lo (2002), we use the daily Sharpe ratio to minimize the estimation error.

D. Computing p-values Using a Randomization Test

This study evaluates the models based on three performance metrics: the directional chi-squared value χ_d^2; the compounded wealth W; and the daily Sharpe ratio SR. The metrics are evaluated over a time period of N steps. The statistical significance of the measured

performance metrics is computed with respect to the null hypothesis

$$H_0 \;=\; \text{the models are random trading strategies with identical trading patterns, drawn so as to respect the signal correlation between the models.}$$

The set $\mathcal{S}$ of all possible strategies, neglecting the trading pattern, is generated by the set of signal types $\mathfrak{s}$ as $\mathcal{S} = \mathfrak{s}^{\times N}$ (were $\times$ denotes the Cartesian product). The set $\mathcal{S}$ is equivalent to the set of all sequences of signals in $\mathfrak{s}$, of length N. In our particular case, $\mathfrak{s} = \{-1,\, 0,\, +1\}$ and $N \sim 20$ years ~ 5000 trading days, amounting to a total number of $\sim |\mathfrak{s}|^N \sim 3^{5000} \sim 10^{2386}$ possible strategies.

We are now going to construct a randomization test, following the framework of Romano and Wolf (2005) and Meinshausen et al. (2011), so as to leave invariant the null distribution resulting from the null hypothesis. The matching of the trading pattern includes the numbers $n_{\mathfrak{s}_i}$ of signals of type $\mathfrak{s}_i$ (i.e. long, short and out of the market), as well as the mean $\mu_{\mathfrak{s}_i}$ and standard deviation $\sigma_{\mathfrak{s}_i}$ of the duration of each type of signal. To randomize the signals of a set of strategies, while respecting the trading pattern and model correlations, let us consider the signal matrix

$$\mathfrak{S} = \begin{bmatrix} s_{1,\,1} & \dots & s_{\mathfrak{M},\,i} \\ \vdots & & \vdots \\ s_{1,\,N} & \dots & s_{\mathfrak{M},\,N} \end{bmatrix}, \tag{11}$$

which contains in row i the signals of the $\mathfrak{M} = |M|$ models at time i, and in column j the N signals of model M_j over time. We define $\mathbf{G}(N)$ as the group of permutations of length N. An element $g \in \mathbf{G}(N)$ permutes the rows of the signal matrix as $(g\mathfrak{S})_{i,\,j} = \mathfrak{S}_{g(i),\,j}$. The transformation g preserves the number of signals of each type in each column, and the signal correlation between all column pairs. However, the mean $\mu_{\mathfrak{s}_i}$ and standard deviation $\sigma_{\mathfrak{s}_i}$ of the position durations are not guaranteed to be preserved.

To preserve the position characteristics, let us consider a vector of signals $\vec{s} = (s_1,\, \dots,\, s_N)$. Adjacent signals of the same type can be compressed to a position, which is a pair of a signal and a duration, compressing the sequence to $\vec{s} = \big((s_1,\, l_1),\, \dots,\, \big(s_{n_p},\, l_{n_p}\big)\big)$, where n_p is the number of positions. Permuting the positions does not guarantee that adjacent positions are of different types. When adjacent positions have the same signal, they can be added as $(s,\, l_1) + (s,\, l_2) = (s,\, l_1 + l_2)$, leading to a position with longer duration. Therefore, not all permutations preserve the position durations. To generate permutations that preserve position durations, let us recall that any permutation can be decomposed into a sequence of pairwise permutations. This property can be used, as described in Algorithm 1, to sample randomly a permutation g in the set of permutations that satisfy $\forall i \in \{1,\, \dots,\, n_p - 1\}\;\; (g\vec{s})_i\,[0] \neq (g\vec{s})_{i+1}\,[0]$, where $(\vec{s})_i\,[0] = (s_i,\, l_i)\,[0] = s_i$.

A permutation $g' \in \mathbf{G}(n_p)$ generated by Algorithm 1 can be mapped to a permutation $g' \to g \in \mathbf{G}(N)$ that permutes block-wise the adjacent signals of each type, each block

Algorithm 1 Randomizing a sequence of trading positions. This algorithm finds pairwise permutations of positions, such that after applying the permutation every pair of adjacent positions is of different signal type. These pairwise permutations are multiplied sequentially, until the correlation to the initial signal sequence is below the threshold τ. The correlation function *corr* computes the correlation of the uncompressed signal sequences.

Ensure: $\tau \leq 0,\ \tau << 1$
Input: $\vec{s} = \left((s_1,\, l_1)\,,\, \ldots,\, \left(s_{n_p},\, l_{n_p}\right)\right),\ \tau$
  $\vec{s'} = \text{copy}(\vec{s})$
  $g = (1,\, \ldots,\, n_p) \in \mathbf{G}\,(n_p)$
  while $corr\left(\vec{s},\, \vec{s'}\right) > \tau$ **do**
    $(i,\, j)$:= random position pair in $\{1,\, \ldots,\, n_p\}^2$
    if $s_{i-1} \neq s_j$ and $s_{i+1} \neq s_j$ and $s_{j-1} \neq s_i$ and $s_{j+1} \neq s_i$ **then**
      $s'_i,\, s'_j = s'_j,\, s'_i$ (exchange components i and j)
      $g_i,\, g_j = g_j,\, g_i$
    end if
  end while
  return g

corresponding to a position. Applying the permutation g to the signal matrix $\mathfrak{S}$ of a single strategy ($\mathfrak{S} \in \mathfrak{s}^{\times(N,\,1)}$) preserves the number of positions of each type, and their duration distribution. However, when permuting a signal matrix of multiple strategies, only the trading pattern of the strategy for which the permutation has been constructed is preserved. Nonetheless, permuting signals with a block-size equal or greater then the mean position duration preserves at first order the mean position duration. Therefore, for a signal matrix with strategies of identical trading pattern, applying a permutation generated based on one of the strategies preserves at first order the duration distribution of all strategies.

Our algorithm to generate transformations of a signal matrix does not perfectly satisfy the requirement of the null hypothesis. Nonetheless, for practical purposes, the null distribution can be considered invariant under the transformations $g \in \mathbf{G}\,(N)$ generated by our algorithm. The only necessary precaution to take, in order not to violate the familiwise error rate, is to randomize with respect to the strategy whose trading pattern produces the distribution of the performance metric least prone to reject the null hypothesis. To this end, the Appendix B provides a simulation study of the wealth quantiles p_W as a function of the mean position duration. The strategies with the shortest duration have the distribution least prone to reject the null hypothesis for the compounded wealth metric. Therefore, the randomization of a signal matrix has to be performed with respect to the strategy in the matrix with shortest mean position duration.

The p-values for the null hypotheses of the $\mathfrak{M}$ models, with signal matrix $\mathfrak{S}$, can now be

computed for a performance metric $h : \mathfrak{s}^{\times(N,\mathfrak{M})} \to \mathbb{R}^{\mathfrak{M}}$, following the algorithm of Romano and Wolf (2016). The performance metric maps the signal matrix to a vector of $\mathfrak{M}$ values. The Algorithm 1 is used to generate a set $\mathfrak{g} = \{g_1, \ldots, g_K\}$ of K transformations, which sample the distribution of the performance h under the null hypothesis as $h^*_\kappa = h\left(g_\kappa \mathfrak{S}\right) = \left(h^*_{\kappa,1}, \ldots, h^*_{\kappa,\mathfrak{M}}\right)$, where $\kappa \in \{1, \ldots, K\}$. Then we compute the individual p-values

$$\hat{p}^h_i = \frac{\#_\kappa \left\{h^*_{\kappa,i} \geq h_i\right\}}{K+1}, \tag{12}$$

unadjusted for multiple testing, where h_i is the performance metric for model i, and $i \in \{1, \ldots, \mathfrak{M}\}$. This expression evaluates for model i the number of randomized signals with a performance metric $h^*_{\kappa,i}$ exceeding the actual value h_i of the model. These individual p-values are used to compute the permutation $g_\leq = \{i_1, \ldots, i_\mathfrak{M}\}$, satisfying $\hat{p}^h_{i_1} \leq \hat{p}^h_{i_2} \leq \ldots \leq \hat{p}^h_{i_\mathfrak{M}}$. Finally, the adjusted p-values can be computed from the values

$$\max^*_{\kappa,j} = \max\left\{h^*_{\kappa,i_j}, \ldots, h^*_{\kappa,i_\mathfrak{M}}\right\}, \quad j \in \{1, \ldots, \mathfrak{M}\}, \tag{13}$$

as

$$\hat{p}^{h,adj}_{i_j} = \max\left\{\frac{\#_\kappa \left\{\max^*_{\kappa,j} \geq h_{i_j}\right\}}{K+1}, \hat{p}^{h,adj}_{i_{j-1}}\right\}, \tag{14}$$

where $\hat{p}^{h,adj}_{i_0} = 0$ is used to initialize at i_0 the recursion relation defined by Equation 14.

The error on $\hat{p}^h$ is proportional to $\frac{1}{p^h}$, and accurately estimating $p^h \leq 10^{-6}$ would require to generate millions of randomizations. To speed up the estimation of $\hat{p}^h$ when the value h is far in the tail of its null distribution, we fit the distribution of $\max^*_{\kappa,j}$ on a sample $K = 1000$, and extrapolate the tail probability $\hat{p}^{h,adj}_{i_j}$. The logarithm of compounded wealth and the daily Sharpe ratio are normally distributed according to the central limit theorem. To accommodate for departure from normality, in particular fat tails, we fit a generalized normal distribution (Varanasi, 1989). The directional accuracy is extrapolated using a chi-squared distribution fit.

E. Directional Accuracy & Compounded Wealth

To understand how directional accuracy impacts compounded wealth, we compute an analytic expression of the expected wealth as a function of the trading pattern, the contingency matrix, and the transaction cost. The number of observations of a real return of category i and a predicted return of category j is denoted by $O_{i,j}$. The total number of observations is $N = \sum_{i,j} O_{i,j}$. The directional accuracy is assumed constant in time and independent of the return amplitude. In a first step, we approximate the logarithm of the compounded

wealth as

$$\log(W_N) = \sum_{t=1}^{N-1} \log(1 - \Delta s_t \Delta\varsigma) + \log\left(1 + s_t \cdot r_t + (1 - |s_t|) \cdot r_t^f\right) \quad (15)$$

$$\sim \sum_{i,j} n_{\Delta\mathfrak{s}_{i,j}} \cdot \log(1 - \Delta\mathfrak{s}_{i,j}\Delta\varsigma) \quad (16)$$

$$+ \sum_{i,j} O_{i,j} \cdot \left(\log(1 + \mathfrak{s}_j \cdot \overline{r}_i) + \log\left(1 + \frac{(1 - |\mathfrak{s}_j|)}{1 + \mathfrak{s}_j \cdot \overline{r}_i} \cdot r^f\right)\right). \quad (17)$$

We denote by $\mathfrak{s}_i$ the signal of category i in the contingency matrix, by $\Delta\mathfrak{s}_{i,j} = |\mathfrak{s}_j - \mathfrak{s}_i|$ the absolute difference between two signal categories, and by $n_{\Delta\mathfrak{s}_{i,j}}$ the number of transitions between a signal of type i and a signal of type j. The average asset return of category i is denoted by $\overline{r}_i$, and the average risk free rate by r^f. Introducing the excess predictability

$$\Delta\rho_{i,j} = \frac{1}{N}(O_{i,j} - E_{i,j}), \quad (18)$$

the expression simplifies to

$$\log(W_T) \sim W_{\Delta\varsigma} + W_f + I_\Sigma^p \cdot r_m + N \sum_{i,j} \Delta\rho_{i,j} \log(1 + \mathfrak{s}_j \overline{r}_i), \quad (19)$$

where

$$W_{\Delta\varsigma} = \sum_{i,j} n_{\Delta_{i,j}} \cdot \log(1 - \Delta\mathfrak{s}_{i,j}\Delta\varsigma) \quad (20)$$

is the transaction cost, and

$$W_f = \sum_{i,j} O_{i,j} \log\left(1 + \frac{(1 - |\mathfrak{s}_j|)}{1 + \mathfrak{s}_j \cdot \overline{r}_i} \cdot r^f\right) \quad (21)$$

is the risk free wealth. The third term is the product of the average position and the average market return defined as

$$I_\Sigma^p = N \sum_j \rho_{\cdot j} \mathfrak{s}_j, \quad \text{respectively} \quad r_m = \frac{1}{n_c} \sum_{i,j} \rho_{i\cdot} \log(1 + \mathfrak{s}_j \cdot \overline{r}_i), \quad (22)$$

where the values $\rho_{i\cdot}$ and $\rho_{\cdot j}$ have been defined in equation 7 as $E_{i,j} = N\rho_{i\cdot}\rho_{\cdot j}$. Having more, or stronger, long positions then short positions ($I_\Sigma^p > 0$) in a bull market ($r_m > 0$) will always result in positive returns. The opposite holds true for more, or stronger, short positions in a bearish market.

The fourth term $\sum_{i,j} \Delta\rho_{i,j} \log(1 + \mathfrak{s}_j \overline{r}_i)$ expresses the return stemming from the excess predictability. This term is constrained by the equalities $\sum_i \Delta\rho_{i,j} = 0$ and $\sum_j \Delta\rho_{i,j} = 0$.

These equalities follow from the definitions of $\rho_{i.}$ and $\rho_{.j}$ given in equation 7, which imply $\sum_i \rho_{i.} = 1$ and $\sum_j \rho_{.j} = 1$. Therefore, to further simplify this term in the case of an equal number $n = n_r = n_c$ of real and predicted categories, the excess predictability is assumed to take the homogeneous form

$$\Delta\rho_{i,\,j} = \left(\frac{n}{n-1}\delta_{i,\,j} - \frac{1}{n-1}\right)\Delta\rho, \tag{23}$$

which satisfies the probability constraints $\sum_i \Delta\rho_{i,\,j} = 0$ and $\sum_j \Delta\rho_{i,\,j} = 0$. The excess predictability $n\Delta\rho/(n-1)$ in all correct predictions is homogeneously balanced by the reduced predictability $\Delta\rho/(n-1)$ in all wrong predictions. It has to be noted that in the binary case $n = 2$ this is the only possible expression for $\Delta\rho_{i,\,j}$, while for $n > 2$ less homogeneous excess predictability can arise. The assumption of homogeneous excess predictability transforms the last term into

$$N\frac{n}{n-1}\Delta\rho\sum_i \log\left(1 + \mathfrak{s}_i\overline{r}_i\right) - \frac{N}{n-1}\Delta\rho\sum_{i,\,j}\log\left(1 + \mathfrak{s}_j\overline{r}_i\right). \tag{24}$$

In the case of balanced signals, satisfying $\forall\mathfrak{s}_j \in \mathfrak{s},\ -\mathfrak{s}_j \in \mathfrak{s}$, the second term can be neglected up to second order in $\overline{r}_i$. Finally, introducing the average return $r_c = \frac{1}{n}\sum_i \log\left(1 + \mathfrak{s}_i\overline{r}_i\right)$ of a correct prediction, the compounded wealth can be expressed as

$$\log\left(W_N\right) \sim W_{\Delta\varsigma} + W_f + I_\Sigma^p \cdot r_m + N\frac{n^2}{n-1}\Delta\rho \cdot r_c. \tag{25}$$

The last term in equation 25 is the expected profit made from the excess predictability, which is proportional to the average profit r_c of a correct prediction and the excess predictability $\Delta\rho$. This result is rather intuitive, as making $\Delta\rho$ percent more correct than wrong predictions, with a gain or loss of r_c on each prediction, must yield an overall gain proportional to $\Delta\rho \cdot r_c$. If the compounded wealth differs significantly from the value obtained in equation 25, the directional accuracy is not independent of the return amplitude, as assumed in this computation.

To finally compare the transaction costs and excess profits, let us consider a random trading strategy going only long or short ($n = 2$). A random sequence of binary signals, with identical probability, trades on average every second day. Consequently, the first term evaluates to $\frac{N}{2}\log\left(1 - 2\cdot\Delta\varsigma\right)$, as $\Delta\mathfrak{s}_{i,\,j} = 2\delta_{i,\,j}$ and $n_{\Delta\mathfrak{s}_{i,\,j}} = \frac{N}{4}$. The risk free wealth is zero ($W_f = 0$), as the random strategy is always long or short. The average return of a correct prediction is equal to the average absolute market return ($r_c = \langle\log\left(1 + |r|\right)\rangle$). Therefore, in the case of a flat market ($r_m = 0$), the profits from excess directional accuracy exceed the trading costs if

$$-\frac{1}{2}\log\left(1 - 2\cdot\Delta\varsigma\right) + 4\Delta\rho \cdot r_c > 0 \;\Rightarrow\; \Delta\rho \gtrsim \frac{\Delta\varsigma}{4r_c}. \tag{26}$$

A typical equity index, such as the S&P 500, has an average absolute daily return $\langle |r| \rangle \sim 1\%$. This implies that for every basis point in spread between buy and sell (spread $= 2 \cdot \Delta\varsigma$) at least $100 \times \frac{0.005\%}{4 \times 1\%} = 0.125\%$ in excess predictability $\Delta\rho$ is needed. Using Equation 6, the excess predictability $\Delta\rho$ can be linked to the directional Chi-squared value as $\chi_d^2 = 16N\Delta\rho^2$, assuming $n = 2$ and $E_{i,j} \sim \frac{N}{4}$, $\forall i,\, j \in \{1,\, 2\}$.

III. Empirical Results

We test the universe of models and search technology to forecast daily returns of equity indices in three geographical areas: Asia, Europe and U.S. The selected stock indices are the CSI 300 Index in Shanghai, the FTSE in London and the S&P 500 in the U.S. The daily returns are loaded from the Thomson Reuters data stream, while the monthly risk free rates are taken from the Ken French's data library French (2012). An overview with relevant key values is provided in Table III. The starting date of the CSI 300 correspond to the first ever trading day of that index.

Index	Start	End	W_y(%)	SR_y	↑(%)	r_m (%)	r_c (%)
CSI 300	Apr. 8, 2005	Dec. 31, 2015	13.12	0.45	54.1	$5.60_{\cdot 10^{-3}}$	1.34
FTSE	Jan. 1, 1995	Dec. 31, 2015	3.45	0.11	52.2	$0.88_{\cdot 10^{-3}}$	0.82
S&P 500	Jan. 1, 1995	Dec. 31, 2015	7.37	0.31	53.7	$2.64_{\cdot 10^{-3}}$	0.82

Table III: **Summary statistics of the equity indices used for the empirical study.** This table provides a summary statistics of the equity indices used to test the statistical learning methods on daily returns. The key values in order are: the compounded annual growth rate W_y; the yearly Sharpe ratio SR_y; the number of positive days; the average daily market return r_m due to the market trend as defined in Equation 22; and the average daily return r_c on a winning trade as defined in Equation 25.

A. Best Models

Table IV presents for each model family the statistical significance of the best model. The performance measures are the directional accuracy, compounded wealth, and Sharpe ratio as defined in Section II. The test statistics are computed using a one-sided test for high performance. Cases with significantly low performance were not found. The p-values have been adjusted for multiple testing within each family using the algorithm of Section II.D. These p-values allow us to determine if a model family is significant when considered in isolation.

The p-values p_W of the compounded wealth and p_{SR} of the Sharpe ratio are mostly identical. This finding is not surprising as both performance measures depend primarily on

	S&P 500				FTSE				CSI 300			
Family	$p^{\min}_{\chi^2_d}$	$p^{\min}_{W}$	$p^{\min}_{SR}$	$\#^{95}_{W}$	$p^{\min}_{\chi^2_d}$	$p^{\min}_{W}$	$p^{\min}_{SR}$	$\#^{95}_{W}$	$p^{\min}_{\chi^2_d}$	$p^{\min}_{W}$	$p^{\min}_{SR}$	$\#^{95}_{W}$
GARCH	$\mathbf{6.4}_{\cdot 10^{-3}}$	0.24			0.19	0.67	0.68		$\mathbf{1.3}_{\cdot 10^{-4}}$	0.05		
NN-R	0.84	0.25			0.10	$\mathbf{9.1}_{\cdot 10^{-3}}$		**163**	$\mathbf{2.0}_{\cdot 10^{-2}}$	0.35		
SVM-R	0.11	0.25			$\mathbf{2.2}_{\cdot 10^{-3}}$	$\mathbf{2.2}_{\cdot 10^{-3}}$		**187**	$\mathbf{6.4}_{\cdot 10^{-4}}$	0.35		
DT-R	0.78	0.48			$\mathbf{3.9}_{\cdot 10^{-2}}$	$\mathbf{3.6}_{\cdot 10^{-2}}$		**1**	0.22	$\mathbf{2.1}_{\cdot 10^{-2}}$		**1**
RF-R	0.47	0.55			0.67	0.11			0.43	0.15		
GB-R	0.22	0.20			0.33	0.07			0.16	$\mathbf{1.0}_{\cdot 10^{-2}}$		**1**
LDA-C	0.14	0.26	0.23		$\mathbf{3.3}_{\cdot 10^{-2}}$	$\mathbf{3.7}_{\cdot 10^{-3}}$	$\mathbf{3.4}_{\cdot 10^{-3}}$	**186**	$\mathbf{3.0}_{\cdot 10^{-2}}$	0.56	0.57	
QDA-C	0.41	0.66	0.65		$\mathbf{3.6}_{\cdot 10^{-2}}$	0.06	0.05		$\mathbf{4.4}_{\cdot 10^{-3}}$	0.17	0.18	
LR-C	$\mathbf{4.4}_{\cdot 10^{-3}}$	0.37	0.28		0.19	0.68	0.67		$\mathbf{1.1}_{\cdot 10^{-4}}$	$\mathbf{4.7}_{\cdot 10^{-2}}$	$\mathbf{6.8}_{\cdot 10^{-2}}$	**1**
NN-C	0.78	0.98			0.84	0.10			$\mathbf{9.0}_{\cdot 10^{-6}}$	$\mathbf{2.0}_{\cdot 10^{-2}}$	$\mathbf{1.9}_{\cdot 10^{-2}}$	**150**
SVM-C	$\mathbf{8.7}_{\cdot 10^{-4}}$	0.04	0.07		0.22	0.25	0.36		$\mathbf{3.0}_{\cdot 10^{-4}}$	0.07		
DT-C	0.80	0.78	0.77		$\mathbf{7.3}_{\cdot 10^{-3}}$	0.24			0.86	0.44		
RF-C	0.80	0.89	0.87		0.65	0.28			0.07	0.18	0.16	
GB-C	0.73	0.92	0.91		0.58	0.05			0.11	0.21		
DT-BR	$\mathbf{2.9}_{\cdot 10^{-4}}$	$\mathbf{3.6}_{\cdot 10^{-3}}$	$\mathbf{5.6}_{\cdot 10^{-3}}$	**166**	$\mathbf{7.5}_{\cdot 10^{-3}}$	$\mathbf{2.8}_{\cdot 10^{-3}}$		**196**	$\mathbf{3.5}_{\cdot 10^{-3}}$	0.45		
DT-BC	$\mathbf{7.0}_{\cdot 10^{-6}}$	$\mathbf{6.4}_{\cdot 10^{-4}}$	$\mathbf{9.3}_{\cdot 10^{-4}}$	**119**	$\mathbf{1.6}_{\cdot 10^{-2}}$	$\mathbf{5.9}_{\cdot 10^{-3}}$	$\mathbf{5.5}_{\cdot 10^{-3}}$	**196**	0.14	0.12		

Table IV: **Summary statistics by model family: out-of-sample S&P 500, FTSE and CSI 300.** This table provides for each model family the three best p-values for the directional accuracy ($p_{\chi^2_d}$), the compounded wealth (p_W), and the Sharpe ratio (p_{SR}). The Sharpe ratio p-value is only indicated if $p_{SR} \neq p_W$. The model families are **N**earest **N**eighbors (NN), **S**upport **V**ector **M**achine (SVM), **D**ecision **T**ree (DT), **R**andom **F**orest (RF), **G**radient **B**oosting (GB), **L**inear **D**iscriminant **A**nalysis (LDA), **Q**uadratic **D**iscriminant **A**nalysis (QDA), and **L**ogistic **R**egression (LR). The family names are suffixed as follows: "-R" for regressors; "-C" for classifiers; and "-B" for binary. The p-values are adjusted for multiple testing within each family according to the algorithm in Section II.D. Each family has $49 \times 4 = 196$ models, parametrized by 49 in-sample lengths $l_{is} \in [20, 500]$ and 4 lags $m \in [2, 3, 4, 5]$. The value $\#^{95}_W$ indicates the number of models significant at the 95% level in compounded wealth.

the mean daily return. Exceptions to the identical wealth and Sharpe ratio significance are found among the top performing models. For example, for the best DT-B model on the S&P 500 the large upside volatility, seen on Figure 7, penalizes the Sharpe ratio performance. On the contrary, for the best LDA-C model on the FTSE, staying out of the market during high volatility days improves the Sharpe ratio performance.

In this study, we focus on models with statistically significant compounded wealth, the measure of profits used to test the EMH. Some model families stand out as significant, far above the 95% confidence level. On the S&P 500, the decision tree models with binary inputs (DT-BR and DT-B) are highly significant. On the FTSE, the best model is the support vector regression (SVM-R), second is the decision tree regressor with binary inputs (DT-BR), and third comes the linear discriminant analysis (LDA-C). Further significant models are the nearest neighbor regression (NN-R) and binary decision trees (DT-BC). The single case above the 95% level for the decision regressor model family (DT-R) could be a statistical outlier. On the CSI 300, the nearest neighbor classifier model family (NN-C) is consistently significant. However, it is significantly outperformed by a single GB-R model.

To better understand the performance of the models in the top performing family for each equity index, their compounded wealth has been plotted against their directional accuracy. Figure 2 shows the DT-BC models on the S&P 500. The out-performers are found at lags two and three, with lag three visibly out-performing lag two. Figure 3 shows the SVM-R models on the FTSE. None of the lags is significantly out-performing the others, with the exception of one visible outlier at lag three. Figure 4 shows the NN-C models on the CSI 300, with significant out-performers at lag three and four.

For classification models, the compounded wealth should follow closely the expected value $E\left(W\left(p_{\chi^2_d}\right)\right)$ of Equation 25, as the binary outputs used for training do not contain any information about the return amplitude. Indeed, for the S&P 500 and the CSI 300, the classification models follow quite closely the theoretical relation, up to a constant shift. On the S&P 500, the models actual wealth outperforms the expected wealth consistently, indicating that predictability was higher during volatile periods. On the CSI 300, the models actual wealth are in good agreement with the expected wealth for lag four, while being slightly lower for lag three. This indicates that the lag three predictors have missed some major tail events.

For regression models, the compounded wealth is not necessarily linked to overall directional accuracy, as the prediction can be strongly dependent on the return amplitude. Indeed, for the top performing regression model on the FTSE (SVM-R), there is no strong relationship between directional accuracy and compounded wealth. The models merely lie within the shape of the confidence intervals computed from random strategies.

The confidence regions of randomized strategies on the CSI 300 are strongly stretched along the W axis, which is in part due to the fact that the time period 2007-2015 is less then half the one used for the FTSE and S&P 500. This stretched confidence region is a result

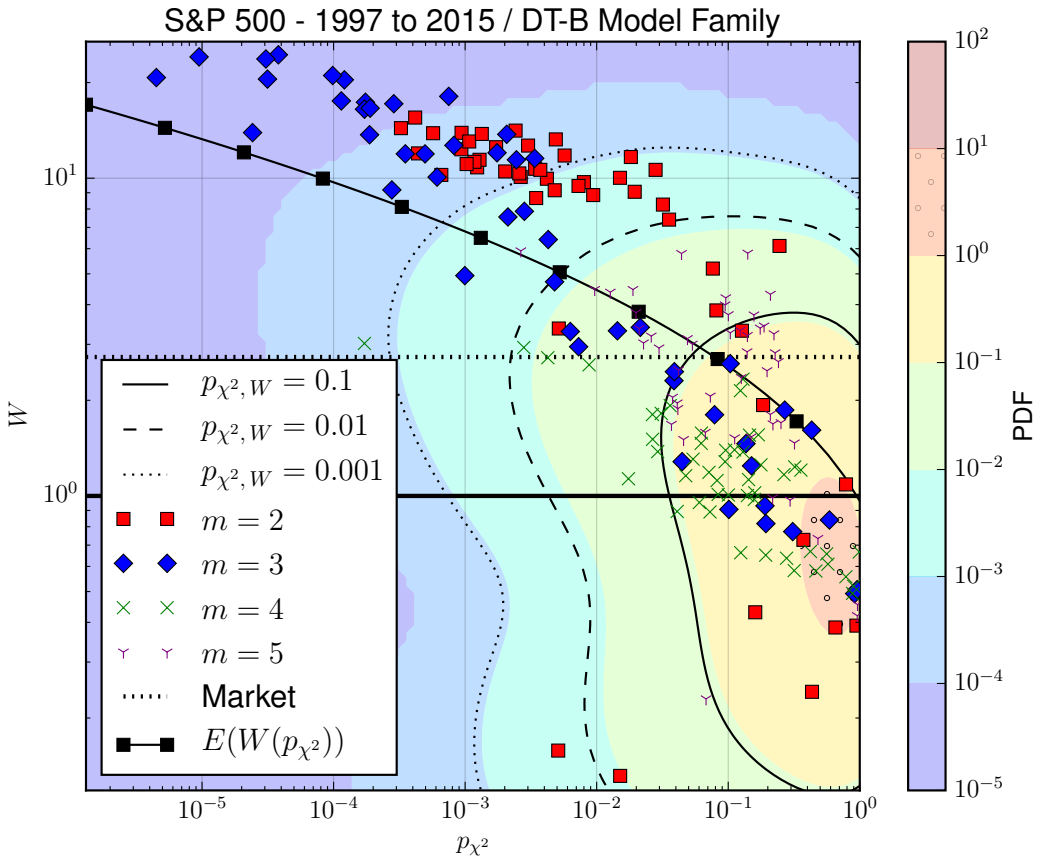

Figure 2: **The compounded wealth (W) versus the directional accuracy p-value ($p_{\chi^2_d}$) for the binary decision tree model family (DT-B), on the S&P 500 between 1997 and 2015.** The trading days between 1995 and 1997 are only used for calibration, but not for the out-of-sample performance. The four different lags $m \in [2, 3, 4, 5]$ are differentiated by distinct markers. The market return during the period is given as a reference by a horizontal dotted line. The probability distribution function, as well as the confidence regions $p_{\chi^2_d,W}$, are obtained using one million simulated random strategies. The expected wealth $E\left(W\left(p_{\chi^2_d}\right)\right)$ as a function of the directional accuracy p-value is computed by Equation 25, at zero transaction cost $\Delta\varsigma = 0$.

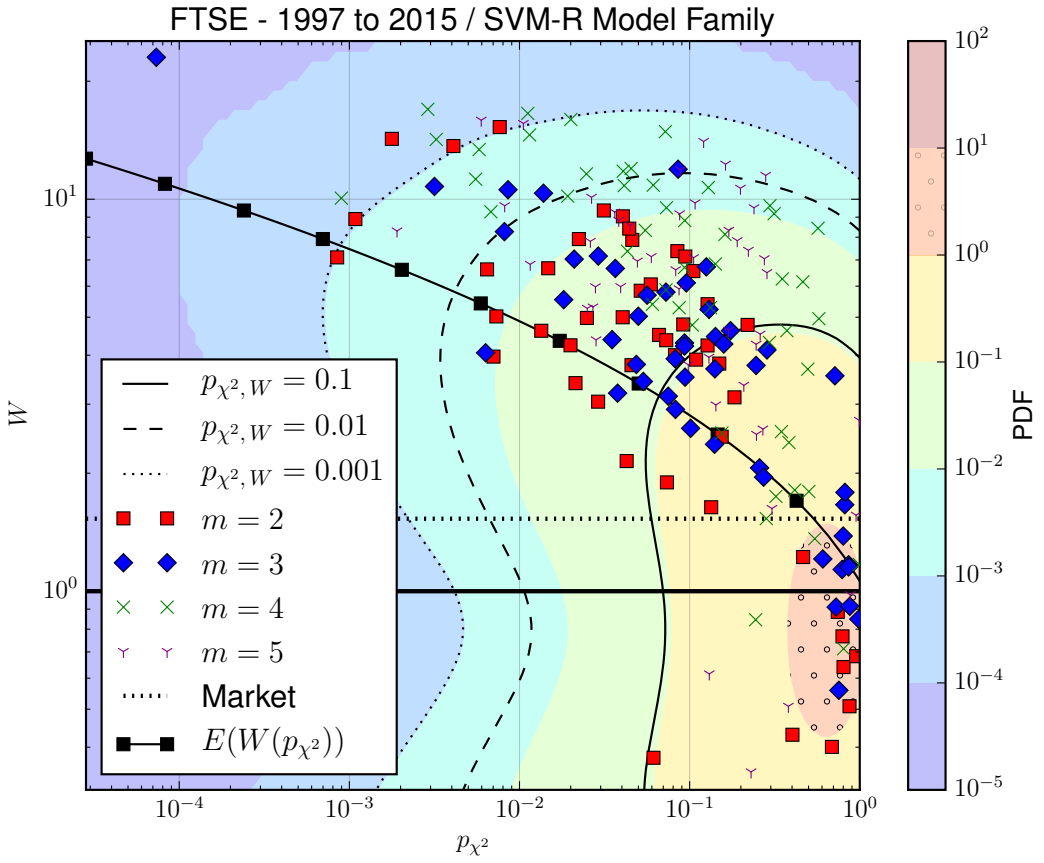

Figure 3: The compounded wealth (W) versus the directional accuracy p-value ($p_{\chi^2_d}$) for the support vector machine regression model family (SVM-R), on the FTSE between 1997 and 2015. For details see caption of Figure 2.

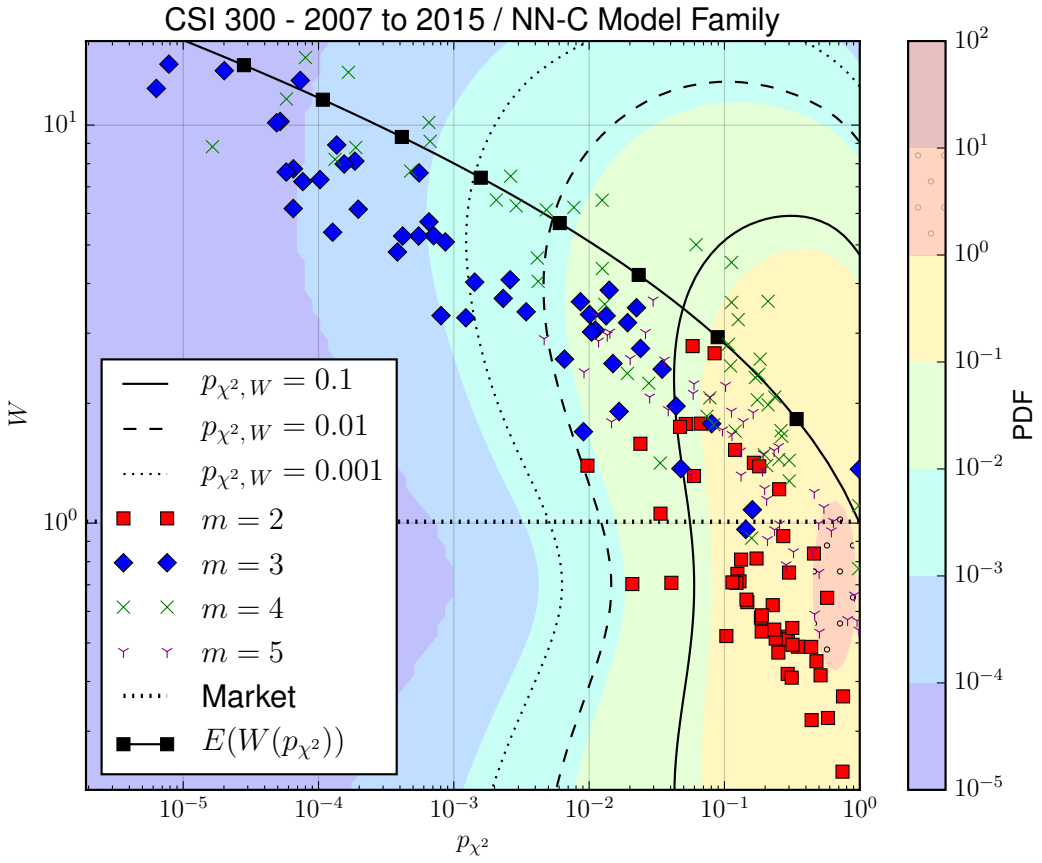

Figure 4: The compounded wealth (W) versus the directional accuracy p-value ($p_{\chi^2_d}$) for the k nearest neighbor classifier model family (NN-C), on the CSI 300 between 2007 and 2015. For details see caption of Figure 2.

of fat tails in the return distribution. Consequently, on the CSI 300, more then half of the randomized strategies lie above the expected wealth curve.

To check for redundancy in our model universe, we compute the Pearson correlation of the model families, averaging all correlations at identical in-sample length and lag. The correlations averaged over all three equity indices are shown in Figure 5. Most model pairs have a correlation below 50%, clearly justifying their individual appearance in the model universe.

High correlations are found among the (SVM-R, SVM-C, LDA-C, LR-C) models that correlate up to 84%. This is expected, as these methods all find a hyperplane maximally separating up and down moves. In case of the LDA-C method, the hyperplane is separating the two multivariate distributions of the two classes. For the LR-C method, the hyperplane is maximizing the asymmetry of up and down moves on each side of the plane. The high correlation between the SVM-R and LDA-C models explains why both families are significant simultaneously on the FTSE (see Table IV). Two other well correlated pairs are (DT-BR, DT-B) and (NN-R, NN-C) with correlation coefficient equal to 68%, respectively 63%. This is expected too, as both models in these two pairs differ only in taking the average or majority vote among an identical group of samples. The high correlation between the DT-BR and DT-B models explains why both families are significant simultaneously on the S&P 500 and FTSE. Last, the LR-C method correlates at 65% with the GARCH model, indicating that it performs a similar type of trend following.

Given the model families that are highly significant in isolation, the next step is to compute the significance over the entire universe of models. The p-values adjusted for multiple testing across the entire model universe are computed using the algorithm from Section II.D. Based on the simulation study of Appendix B, we know that the mean position duration of the randomized strategies needs to be as close as possible to the actual strategy, to maximize the statistical power of the test. However, the mean position duration cannot exceed the actual value, otherwise the familiwise error rate would be violated.

To determine the optimal randomization scheme, we present in Figure 6 the mean duration of all model families. Except for the GARCH, LR-C, and SVM-C models, the mean position duration is always close to two days. Given the low correlation of these three models with the other 13 models, we randomize these two groups of models independently. While sampling independently reduces somewhat the statistical power of the test, this is more then offset by the improvement of removing the three models with long mean position duration. As can be read from Figure 10, longer mean position durations imply significantly lower wealth quantiles, and therefore these strategies do not impact the sampling of the wealth tail of randomized strategies with shorter mean position duration.

An overview of the models with highest significance on p_W is presented in Table V. The p-values adjusted for multiple testing across the entire universe remain significant for the FTSE and S&P 500. The best model for the CSI 300 is significant at the 94% level after only 8 years of trading, and would likely be highly significant on an equivalent 18 year period

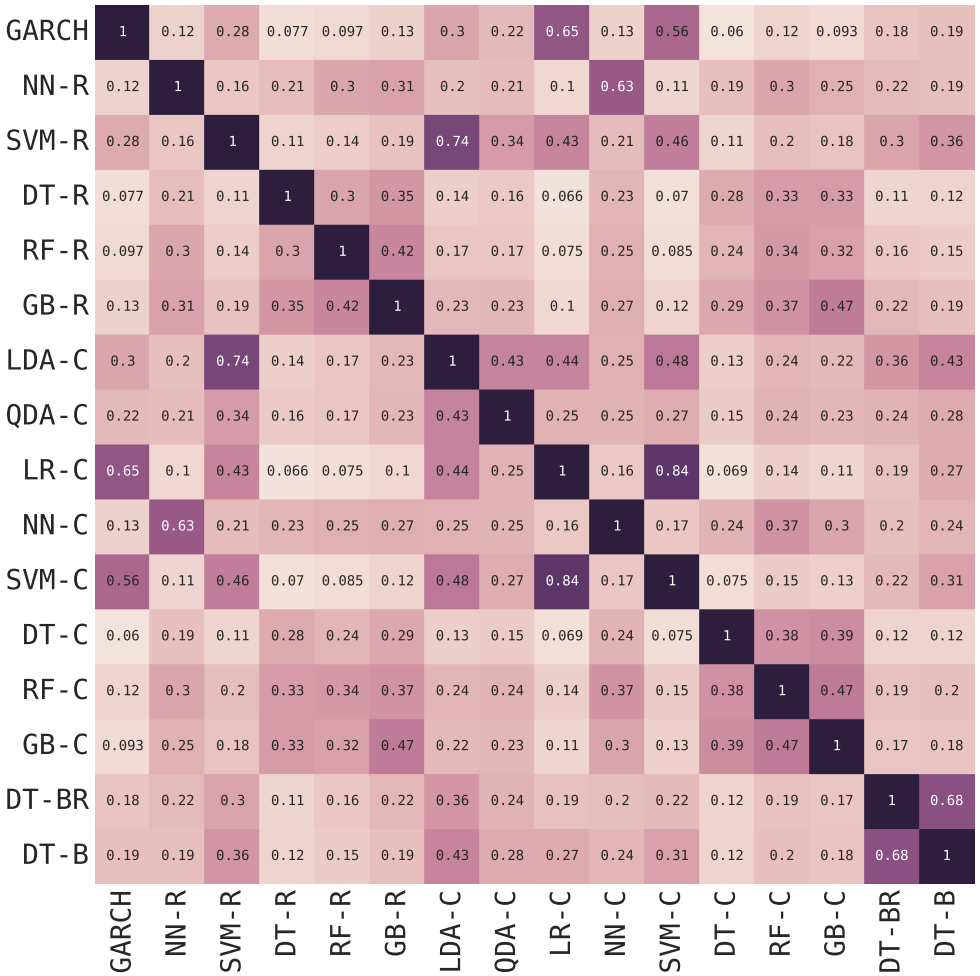

Figure 5: **Pearson correlation of returns between the different model families, at identical in-sample length and lag.** The Pearson correlations have been averaged over in-sample length l_{is}, lags m, and the three equity indices (S&P 500, FTSE, and CSI 300). The models are all positively correlated due to excess in up moves in the predicted asset, inducing a proportional excess in up predictions by the models. Most model pairs have a correlation below 50%, resulting from their fundamentally different prediction dynamics. None of the models pairs are redundant above the 84% level.

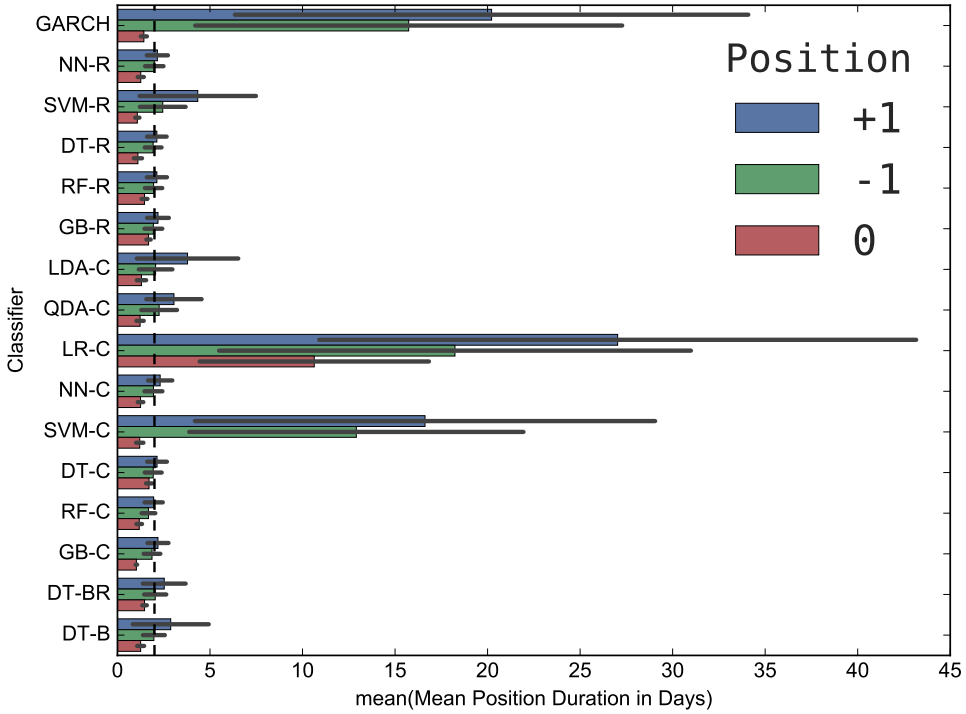

Figure 6: **Mean position duration in days by model family.** The position durations have been averaged over all lags, in-sample lengths, and equity indices. The error bars indicate the average one standard deviation of the durations for a single model. The model families can be split into two clusters: GARCH, LR-C, and SVM-C with high mean durations between 10 and 25 days; and all the other models with mean durations around two days (vertical dashed line).

if its performance remained at the same level. The best performing model for each equity index is consistently found at lag $m = 3$, while the optimal in-sample length is found in two different regions. For the S&P 500 and the FTSE, the optimal in-sample length clusters around $l_{is} \in [390, 400]$, and for the CSI 300 the optimal in-sample length cluster around $l_{is} \in [150, 160]$. For the equity indices tested in this study, these combinations of lag and in-sample length provide the best tradeoff in maximizing sample size and minimizing the risk of calibrating across multiple regimes.

The trading performance over time of a selection of the best models presented in Table V is shown in Figure 7. For the FTSE, the top performing DT-BC and LDA-C models are not shown as they correlate highly with the shown models DT-BR, respectively SVM-R. This Figure reveals that the highest abnormal returns on the S&P 500 and FTSE occur during the dot-com bubble (1997-2003), the financial crisis (2008-2009), and the European debt crisis (2012). The GB-R and NN-C models on the CSI 300 have very different dynamics. The abnormal returns of the GB-R model correlate highly with the financial crisis (2008), and the recent Chinese stock market turbulences (2015). The GB-R model is a variant of boosted decision trees, and this result therefore strengthens the finding that decision tree based models have significant predictability during crises. The abnormal returns of the NN-C model are made during the period 2009 to 2012, not visibly linked to a crisis. Nonetheless, the financial crisis (2008) and recent Chinese stock market turbulences (2015) are smoothened out in comparison to the buy & hold strategy.

Index	**Model**	m	l_{is}	W_y(%)	SR_y	$p_{\chi^2_d}$	p_W	$\Delta\rho$ (%)	$2\Delta\varsigma^{=bh}$ **(bps)**
CSI 300	**GB-R**	**3**	**160**	40.09	0.49	0.33	0.06	1.23	29.7
	NN-C	4	150	36.55	1.08	$1.8_{\cdot 10^{-3}}$	0.09	2.16	28.5
	NN-C	3	380	33.76	1.06	$1.8_{\cdot 10^{-3}}$	0.09	2.47	25.7
FTSE	**SVM-R**	**3**	**390**	17.92	0.75	$7.2_{\cdot 10^{-2}}$	$2.4_{\times 10^{-2}}$	1.36	15.9
	DT-BR	3	250	16.80	0.43	$9.3_{\cdot 10^{-2}}$	$2.6_{\times 10^{-2}}$	1.11	11.2
	KNN-R	2	350	15.20	0.47	0.27	$3.5_{\times 10^{-2}}$	0.99	9.68
S&P 500	**DT-BC**	**3**	**400**	18.55	0.89	$2.0_{\cdot 10^{-3}}$	$1.9_{\times 10^{-2}}$	1.46	12.9
	DT-BR	3	330	17.90	0.74	$2.1_{\cdot 10^{-3}}$	$1.9_{\times 10^{-2}}$	1.51	9.38

Table V: **Summary of the top performing models on the compounded wealth metric (p_W) for each equity index.** The key values in order are: the model family; the number of lags m; the in-sample length l_{is}; the compounded annual growth rate W_y; the yearly Sharpe ratio SR_y; the p-values $p_{\chi^2_d}$ and p_W adjusted for multiple-testing in the entire universe of models; the excess predictability $\Delta\rho$ as defined in Equation 23; and the round trip transaction costs $2\Delta\varsigma^{=bh}$ breaking-even in W_y with the buy & hold strategy.

To rigorously define the periods of abnormal returns, we studied the distribution of the

performance with respect to the buy & hold strategy. The value $\log\left(\frac{S(t)}{BH(t)}\right)$ is used as a scale free measure of performance, were $S(t)$ is the value of the strategy, and $BH\,(t)$ is the value of the buy & hold strategy. The typical period of abnormal returns is around 6 months, and we hence use the performance measure $P_t = \Delta_{6m} \log\left(\frac{S_t}{BH_t}\right)$, were Δ_{6m} is the 6 months differentiator. The distribution of P over time is found to be non-normal, and we therefore define as over-performance any value of P in the upper non-normal tail. Figure 8 shows the performance P over time for the same models as Figure 7, as well as the distribution of the performance P aggregated over the models for each equity index. The periods of excess performance confirm their correlation with the crisis periods mentioned previously.

B. Ex-ante Performance

Our finding of profitable model families, statistically significant after correcting for data snooping, does not guarantee that it would have been possible to select these models ex-ante. To settle this question, we apply the best Sharpe ratio search technology to the entire universe of models, for each equity index. The performance of the resulting portfolios are presented in Table VI. The search technology is only significant for the S&P 500, and fails at selecting the best model for the FTSE and CSI 300.

These results are not surprising given the model family performances of Table IV. The S&P 500 has one model family (DT-BB) that outperforms the second best model family (DT-BR), while all the other families are insignificant. This performance distribution among the models makes it easy for the search technology to select the best model and stick to it. For the FTSE and the CSI 300, multiple model families are competing at similar significance levels. This induces the search technology to constantly switch to the new best performing model. Often the switch occurs at the moment were the best performing model suffers a draw-down, leaving the search technology portfolio with a mediocre performance compared to the best performing model.

The Figure 9 shows the trading performance over time of the search technology portfolio on all three equity indices. On the CSI 300, the search technology keeps switching models until 2011, with an unfavorable timing producing abnormally low returns. Starting in 2011, the search technology sticks to the best performing model of Table V (NN-C, $l_{is} = 380$, $m = 3$), producing abnormally high returns until the end of 2015. On the FTSE, the search technology keeps switching models during the entire time period, never producing any abnormal returns. On the S&P 500, the search technology selects the best model of Table V (DT-BB, $l_{is} = 400$, $m = 3$) early on, and benefits from the constant abnormal performance of this model.

The performance achieved on the S&P 500 remains significant at the 99.7% level even after adjusting for multiple testing on three independent equity indices. The transaction costs breaking even with the market are $2\Delta_\varsigma \sim 10$ bps per round trip. Assuming a typical

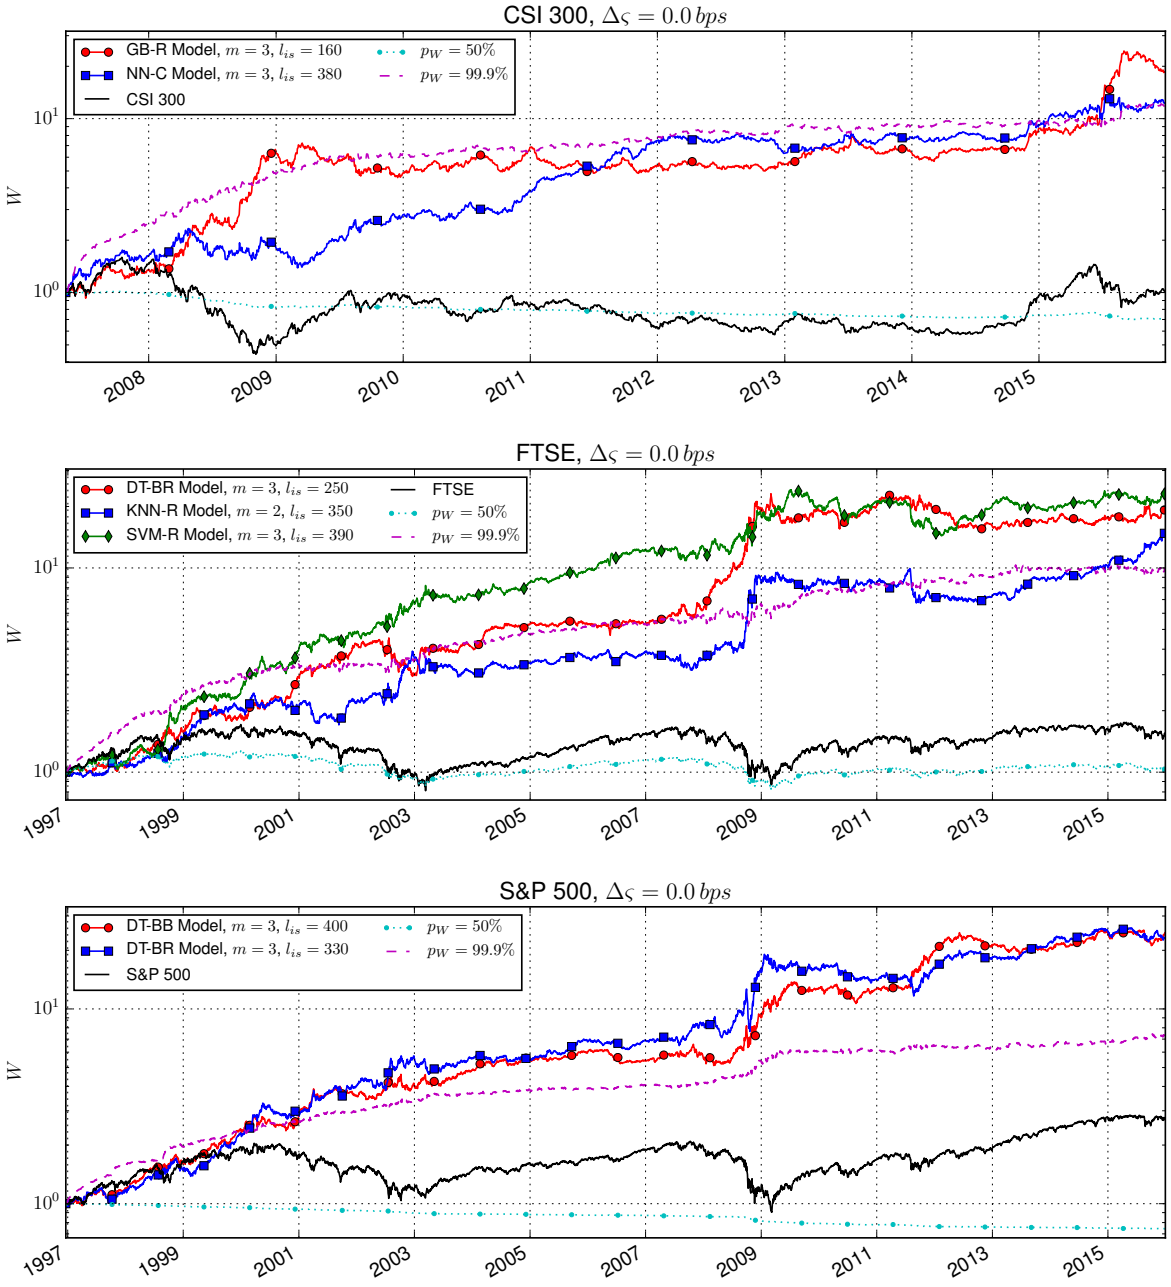

Figure 7: **Trading performance of the best model for each equity index.** These three figures show the compounded wealth $W(t)$ of the best model from Table V for the CSI 300, FTSE, and S&P 500. As reference, the figures show the market return (or buy & hold strategy), as well as the $p_W = 50\%$ and $p_W = 0.1\%$ quantiles of the randomized strategies. All strategies are shown at zero transaction cost $\Delta\varsigma = 0$.

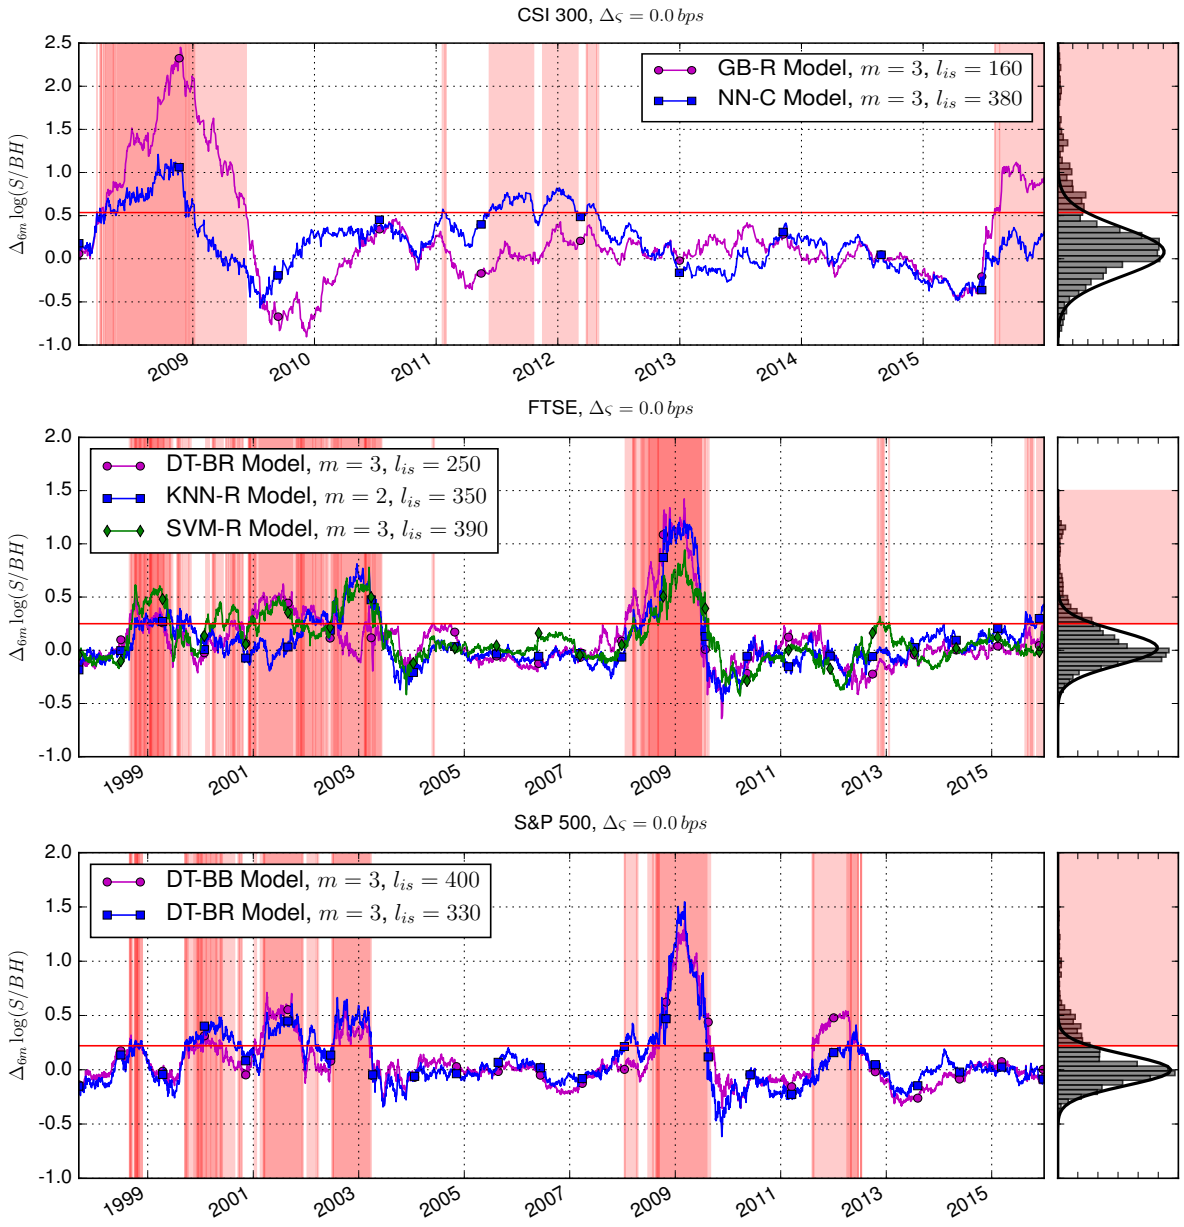
CSI 300, $\Delta_\varsigma = 0.0\,bps$
GB-R Model, $m = 3$, $l_{is} = 160$
NN-C Model, $m = 3$, $l_{is} = 380$
$\Delta_{6m} \log(S/BH)$
FTSE, $\Delta_\varsigma = 0.0\,bps$
DT-BR Model, $m = 3$, $l_{is} = 250$
KNN-R Model, $m = 2$, $l_{is} = 350$
SVM-R Model, $m = 3$, $l_{is} = 390$
S&P 500, $\Delta_\varsigma = 0.0\,bps$
DT-BB Model, $m = 3$, $l_{is} = 400$
DT-BR Model, $m = 3$, $l_{is} = 330$

Figure 8: **Excess trading performance with respect to buy & hold.** The performance P is measured as $P_t = \Delta_{6m} \log\left(\frac{S_t}{BH_t}\right)$, were S_t is the value at time t of the strategy, BH_t is the value of the buy & hold strategy at time t, and Δ_{6m} is the 6 months differentiator. These three figures show the performance measure of the best models (GB: **G**radient **B**oosting; NN: **N**earest **N**eighbor; DT: **D**ecision **T**ree; SVM: **S**upport **V**ector **M**achine; suffixed as "-R" for regressors, "-C" for classifiers, and "-B" for binary) from Table V for the CSI 300, FTSE, and S&P 500. Over-performance is determined by the threshold (red line) were the tail of the performance distribution becomes non-normal.

Index	W_y(%)	SR_y	$p_{\chi^2_d}$	p_W	$\Delta\rho$ (%)	$2\Delta\varsigma^{=bh}$ **(bps)**
CSI 300	-3.41	0.19	$2.5_{\cdot 10^{-2}}$	0.48	1.31	-0.6
FTSE	3.16	0.16	0.40	0.15	0.31	2.2
S&P 500	14.43	0.62	$9.8_{\cdot 10^{-3}}$	$1.0_{\times 10^{-3}}$	0.98	9.9

Table VI: **Summary of the best Sharpe ratio search technology portfolio performance for each equity index.** The key values in order are: the compounded annual growth rate W_y; the yearly Sharpe ratio SR_y; the p-values $p_{\chi^2_d}$ and p_W; the excess predictability $\Delta\rho$ as defined in Equation 23; and the round trip transaction costs $2\Delta\varsigma^{=bh}$ breaking-even in W_y with the buy & hold strategy.

average transaction cost on future of 5 bps per round trip, this result rejects the EMH on the S&P 500.

To determine if this performance can be explained by known factors, we evaluated the performance with the CAPM, the three-factor model of Fama and French (1993), and the four-factor model of Carhart (1997). The full four-factor model measures performance as a time-series regression of

$$R_t - R_{ft} = f_a + f_b\left(R_{Mt} - R_{ft}\right) + f_s SMB_t + + f_h HML_t + f_m MOM_t + e_t. \tag{27}$$

In this regression, R_t is the portfolio return on month t, R_{ft} is the risk-free rate (the 1-month U.S. Treasury bill rate), R_{Mt} is the market return, SMB_t and HML_t are the size and value-growth returns of Fama and French (1993), MOM_t is the momentum return, f_a is the average return not explained by the benchmark model, and e_t the residual error term. The values for R_{ft}, R_{Mt}, SMB_t, HML_t and MOM_t are taken from Ken French's data library (French, 2012), and derive from underlying stock returns data from the Center for Research in Security Prices (CRSP). The three-factor model is obtained by leaving out the momentum term, and the CAPM is obtained by further leaving out the SMB and HML factors.

Table VII shows the intercept and regression slopes (load) for all three models, including their t-statistics. The momentum (MOM), size (SMB), and value growth (HML) slopes are almost zero, and these factors do not explain the observed returns. The market factor correlation with the portfolio is around 10%, but is insignificant ($t\left(R_M\right) = 1.53 < 2$). This correlation is in line with the factor $I^p_\Sigma \cdot r_m$ of Equation 25, expressing the return from the average position I^p_Σ and the market return r_m. As a consequence, the best Sharpe ratio portfolio strategy has a monthly intercept of 0.98%, significant at $t\left(f_a\right) = 2.77$. The intercept remains significant at the threshold $t\left(f_a\right) = 2$ up to transaction costs of $2\Delta\varsigma = 3.6$ bps per round trip. The returns of the best Sharpe ratio portfolio on the S&P 500 can therefore not be explained by known factors, and qualify as a new anomaly.

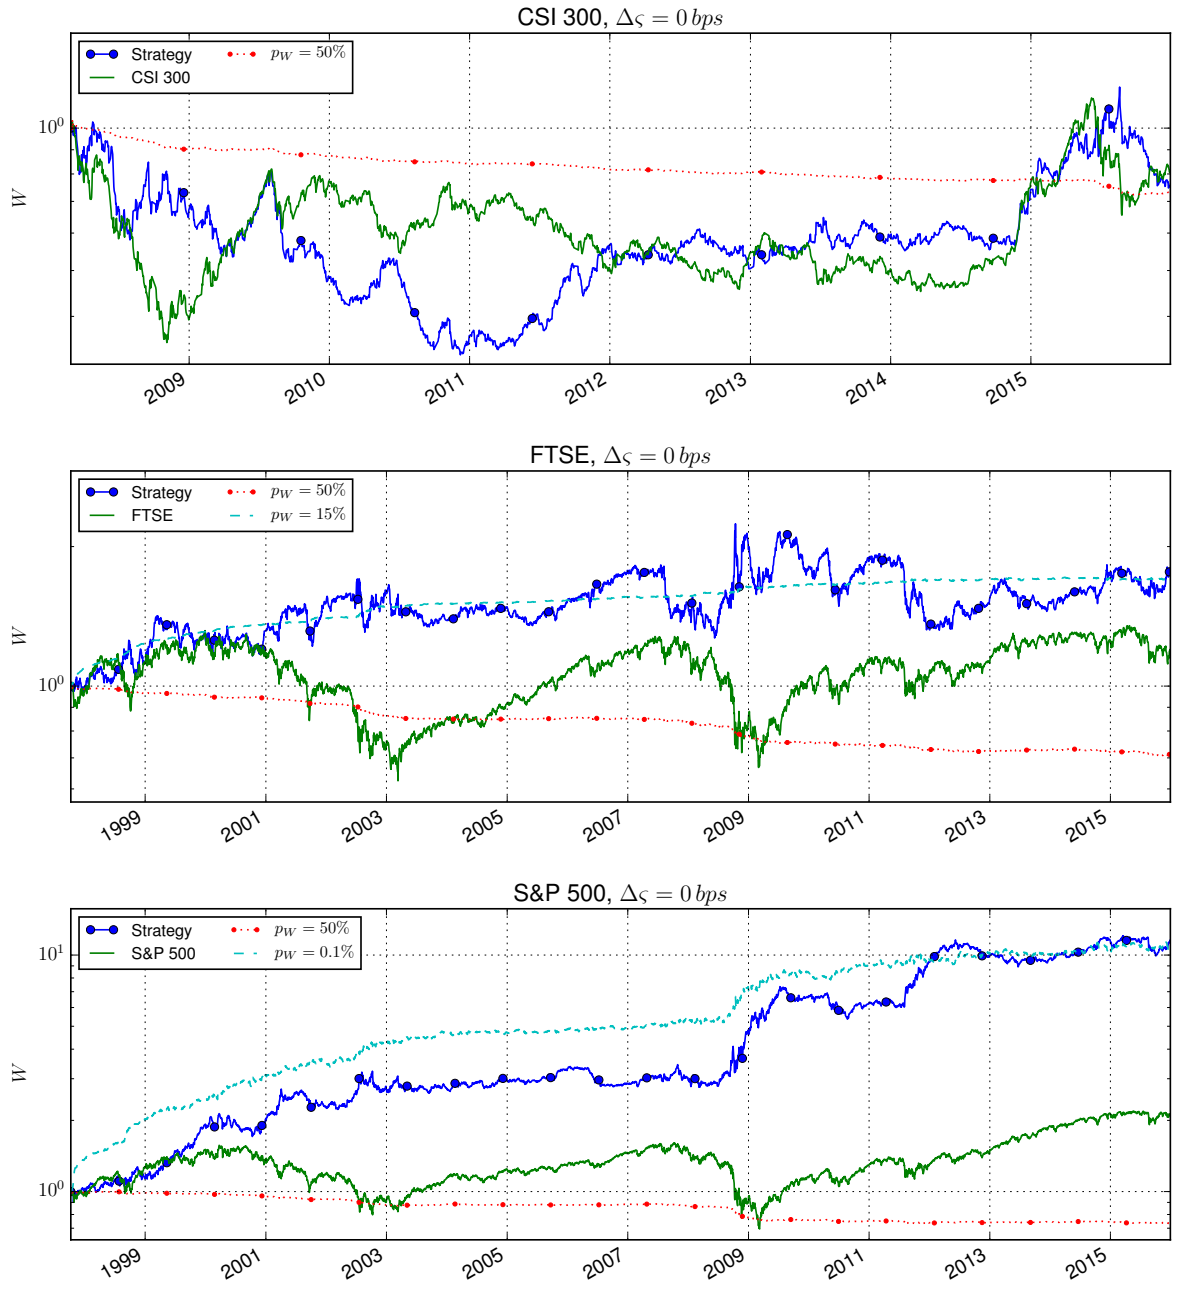

Figure 9: **Trading performance of the search technology portfolio for each equity index.** These three figures show the compounded wealth $W(t)$ of the best Sharpe ratio search technology, applied to the whole universe of models, for the CSI 300, FTSE, and S&P 500. As reference, the figures show the market return (or buy & hold strategy), as well as the $p_W = 50\%$ and break even quantiles of the randomized strategies. All strategies are shown at zero transaction cost $\Delta\varsigma = 0$.

	f_a	f_b	f_s	f_h	f_m	R^2
$Coef$	0.98	0.12				0.05
$t\,(Coef)$	2.77	1.53				
$Coef$	0.98	0.10	0.06	-0.05		0.05
$t\,(Coef)$	2.76	1.24	0.51	-0.47		
$Coef$	1.00	0.09	0.06	-0.06	-0.03	0.05
$t\,(Coef)$	2.78	1.00	0.55	-0.55	-0.39	

Table VII: **Intercepts and slopes in variants of regression for the search technology portfolio on the S&P 500**. The table shows the monthly intercepts (f_a) and regression slopes (f_b, f_s, f_h and f_m, for $R_M - R_f$, SMB, HML, and MOM, respectively), as well as their t-statistics, for the CAPM, three-factor, and four-factor versions of regression. The factors are estimated for the best Sharpe ratio search technology portfolio on the S&P 500 between Jan. 1, 1998 and Dec. 31 2015 at zero transaction cost, as shown in Figure 9. The monthly intercepts are identical within 2% for all versions, and are therefore significant at a t-ratio of $t = 2.76$.

IV. Conclusion and Future Work

We have performed an exhaustive performance analysis of statistical learning methods to forecast daily returns. The model universe was carefully constructed to exhaust the list of all models that can meaningfully be calibrated on the available data. The performance p-values have been adjusted for multiple testing using a variant of the randomization test by Romano and Wolf (2005, 2016). The empirical results of compounded profits of the best models for the S&P 500 and FTSE are found significant above the 97.5% level. The results for the CSI 300 are significant only at the 94% level, nonetheless impressive given the short time period available.

The three best models, one for each equity index, were found at lag $m = 3$, in-sample length $l_{is} \in [390, 400]$ for the S&P 500 and FTSE, and in-sample length $l_{is} = 160$ for the CSI 300. While it could be a statistical fluke, the consistency at lag $m = 3$ strengthens our finding of dependencies in the daily returns of the analyzed equity indices. The results clearly reject the martingale hypothesis $E\,(r_t|r_{t-3},\, r_{t-2},\, r_{t-1}) = 0$, in line with the argumentation of Grossman and Stiglitz (1980) and the Adaptive Market Hypothesis (Lo, 2012). The S&P 500 and FTSE exhibit significant predictability in directional accuracy of the binary return sequences (DT-BR, and DT-B model families), which go unnoticed in linear regression models such as GARCH. This finding is in good agreement with the study by Christoffersen and Diebold (2003). The CSI 300 exhibits significant local predictability among the five nearest neighbors of sequences of three and four returns (NN-C model families), a result that would

go unnoticed as well in linear regression models.

The trading performance over time of the best models was most abnormal during the dot-com bubble (1997-2003), the financial crisis (2008-2009), the European debt crisis (2012), and the recent Chinese stock market turbulences (2015). The EMH seems to hold except during the market crises, where statistically significant deviations are found. During normal periods, the performance of decision tree based strategies is no better than chance. But during crises, pockets of predictability arise, which are discovered by decision tree based methods. While no certain explanation can be given, it would not be surprising to find that behavioral heuristics, as exposed by Tversky and Kahneman (1974) even among trained statisticians, have significant impact during a market crash. Indeed, crisis periods are certainly the most propitious to drastic actions, driven by shortsighted human behavior, and not backed by statistics. As well, in times of fear and panic, investors tend to herd according to the psychology of "being safe in numbers", a behavior reducing the number of competing strategies, and creating a market dominantly driven by a single effective agent (i.e. the herding investors). Coupled with periodic announcements of major monetary institutions, such market dynamics have a high likelihood to introduce systematic biases that can be arbitraged.

To verify the EMH as defined by Timmermann and Granger (2004), we used a best Sharpe ratio search technology. The EMH is rejected at transaction costs below 2.2 bps per round trip for the FTSE, and below 9.9 bps for the S&P 500. Assuming an average cost of 5 bps per round trip for a future on the S&P 500 during that time period, we reject the EMH for the S&P 500. The returns of the best Sharpe ratio portfolio were regressed on the CAPM, three-factor, and four-factor models. In all three regressions, the intercept was significant at $t(f_a) \geq 2$ for the S&P 500 up to transaction costs of 3.6 bps per round trip. Therefore, these anomalous returns cannot be explained by known factors such as the market return or momentum, and constitute a new anomalous 3-day crisis persistency factor, a persistency strongly linked to market crash periods.

Comparing the numerous studies reviewed by Atsalakis and Valavanis (2009), which apply statistical learning methods to financial markets, is challenging given the variety of performance metrics and datasets used. Nonetheless, with the randomization methodology of Romano and Wolf (2005), any result can be given a confidence level that is easy to compare across studies. For example, in our model universe, the highest percentage of correct directional predictions was 52.8% on the 8 year period of the CSI 300, at a confidence level adjusted for data-snooping of $p_{\chi^2_d} = 1.8 \cdot 10^{-3}$. We believe that statistical learning methods, and the non-parametric tests that go along with them, deserve standardized usage alongside the commonly used regression models. Our results show promising applications to the detection of regime changes, which could not be detected at a significant level using unit root tests.

This paper opens up the path to several future research questions. The model universe

could be extended with trinary decision tree models, and neural networks. However, extracting the found dependencies using unsupervised neural networks could reveal challenging. The data used as input could be extended to include more publicly available data such as dividends, risk free rates, or volatility levels. For equity indices showing significant anomalies, such as the S&P 500, the different sectors should be analyzed to determine which stocks drive the predictability. Further leveraging the randomization test, the search technology could be refined to determine the optimal window size on which to select a model, improving upon the expanding window used in this study.

Appendix A. Issues with Cross-Validation on Asset Returns

All statistical learning methods have in common that they find some approximation of the data used for calibration. However, the predictability for asset returns is low and the error term large, which makes it difficult to distinguish between truly predictable patterns and randomly occurring transient patterns. The standard technique to avoid over-fitting is to split the dataset into a **training** and a **test** set. The calibration is then performed in multiple steps. At each step the calibration on the training set is refined and the performance on the test set is computed. When the performance on the test set reaches a plateau, the calibration is considered to be optimal. Further on, cross-validation can be performed by creating multiple pairs of independent (*training*, *test*) sets based on the dataset. Different techniques exist to combine the statistical learner obtained on the different sets.

These cross-validation techniques all rely on the assumption that the samples are independent, which is not the case for the sample constructed from asset returns. These samples are sequential in time and predicting past returns based on future returns makes limited sense. At every time step, a potential test set would necessarily use the last samples in the dataset used for calibration. To be meaningful, such a test set needs to contain a statistically significant number of samples. However, the last samples have the highest causal relation with the predicted out-of-sample output and are therefore needed for calibration. Consequently, any type of cross-validation is likely to affect prediction performance negatively and will not be investigated within the scope of this paper.

Appendix B. Mean Position Duration Impact on Compounded Wealth

When randomizing a signal matrix as defined in Equation 11 that contains strategies with different trading patterns, it is difficult or impossible to preserve the trading pattern of every strategy. Randomizing with respect to the longest position durations would insufficiently randomize the strategies with shorter position durations. On the other hand, randomizing with respect to the shortest position durations would overly randomize strategies with longer position durations. To maximize the statistical power of the randomization test, without

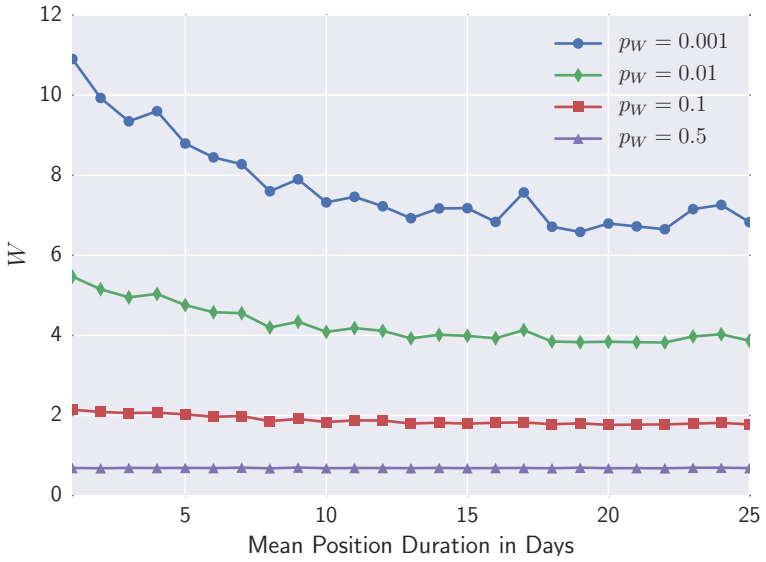

Figure 10: **Compounded wealth quantiles of random trading strategies as a function of the mean position duration in days.** The simulation is based on daily returns of the S&P 500 between Jan. 1, 1997 and Dec. 31, 2015. At each duration in $[1, 2, \ldots, 25]$, the wealth distribution was estimated on a set of 10000 samples. The tail values were evaluated from a generalized normal distribution fitted to the samples. The roughness of the curves is a result of the particular return structure of the S&P 500, and is not a result of an insufficient sample size.

violating the familiwise error rate, we need to compute the dependence of the performance distribution on the position duration.

This dependence has been evaluated for the compounded wealth using a simulation study. Figure 10 shows the wealth quantiles as a function of the mean position duration for the S&P 500, which has the worst case dependence compared to the FTSE and CSI 300. The lower quantiles, below the 95% confidence level, are almost insensitive to the mean position duration. However, larger quantiles decrease significantly with mean position duration. For example, the 99.9% quantile decreases by almost 40% between mean position duration two ($W(2) \sim 11$) and mean position duration 20 ($W(20) \sim 7$). This decrease implies that the randomization has to be performed with respect to the lowest mean position duration in order to ensure the familiwise error rate.

Appendix C. Implementation Assumptions & Precautions

The statistical learning methods presented in section I.D rely on the implementation of the Scikit-learn package from Pedregosa et al. (2011). Scikit-learn is a reliable statistical learning package in Python, which enjoys a good reputation. Mistakes in the statistical learning methods are therefore unlikely, as a large number of people have already approved the good functioning of the Scikit-learn package. Besides the solid implementation of all the methods presented in Table I, this package has an homogeneous interface based on equation 1. This makes it easy to switch methods, as the code can be written at an abstract level using a pointer to an arbitrary statistical learning method.

To minimize the risk of bugs in our code, leading to abnormal predictability, a number of sanity checks were performed. These checks use generated time series for which the expected directional accuracy and average return can be computed analytically. For example, the random walk based on returns $r \in \mathcal{N}(0,\, 1)$ was used to ensure that the directional accuracy and the average return converge to zero, following $\hat{\chi}^2_d \to N^{-1}$ and $\langle r^m_T \rangle \to N^{-\frac{1}{2}}$ as the length N of the time series goes to infinity. Other time series, with controlled predictability based on repeating patterns, were used to verify that the expected directional accuracy and average return are obtained in a variety of scenarios.

Acknowledgments

This study has been made possible by a research grant from ETH Zürich and the interdisciplinary environment of the Chair of Entrepreneurial risks at ETH. In particular, the authors would like to thank D. Ardila, R. Burkholz, Y. Malevergne, H. Hellenes, B. Jónsson, S. Kolbeinsson, B. Vandermarliere, V. Filimonov, Q. Zhang, and D. Daly for their discussions, reviews and inputs.

References

Andersen, J.V., and D. Sornette, 2005, A mechanism for pockets of predictability in complex adaptive systems, *Europhys. Lett.* 70, 697–703.

Atsalakis, George S., and Kimon P. Valavanis, 2009, Surveying stock market forecasting techniques - part II: Soft computing methods, *Expert Systems with Applications* 36, 5932–5941.

Atsalakis, George S., and Kimon P. Valavanis, 2013, Surveying stock market forecasting techniques, part I: Conventional methods.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *The Journal of Finance* 52, 57–82.

Christoffersen, Peter, and Francis Diebold, 2003, Financial asset returns, direction-of-change forecasting, and volatility dynamics .

Fama, Eugene F., 1970, Efficient capital markets: A review of theory and empirical work, *The Journal of Finance* 25, 383–417.

Fama, Eugene F., 1991, Efficient capital markets: II, *The Journal of Finance* 46, 1575–1617.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fernández-Delgado, Manuel, Eva Cernadas, Senén Barro, and Dinani Amorim, 2014, Do we need hundreds of classifiers to solve real world classification problems?, *J. Mach. Learn. Res.* 15, 3133–3181.

French, Ken, 2012, Data library, [Online; accessed 25-July-2016].

Gary, Michael Shayne, and Robert E. Wood, 2010, Mental models, decision rules, and performance heterogeneity, *Strategic Management Journal* 32, 569–594.

Grossman, S.J., and J.E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *The American Economic Review* 70, 293–408.

Hansen, Peter Reinhard, 2005, A test for superior predictive ability, *Journal of Business & Economic Statistics* 23, 365–380.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman, 2001, *The Elements of Statistical Learning*, Springer Series in Statistics (Springer New York Inc., New York, NY, USA).

Hsu, Po-Hsuan, Yu-Chin Hsu, and Chung-Ming Kuan, 2010, Testing the predictive ability of technical analysis using a new stepwise test without data snooping bias, *Journal of Empirical Finance* 17, 471–484.

Hsu, Po-Hsuan, Mark P. Taylor, and Zigan Wang, 2016, Technical trading: Is it still beating the foreign exchange market?, *Journal of International Economics* 102, 188–208.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani, 2014, *An Introduction to Statistical Learning: With Applications in R* (Springer Publishing Company, Incorporated).

Leung, Mark T., Hazem Daouk, and An-Sing Chen, 2000, Forecasting stock indices: a comparison of classification and level estimation models, *International Journal of Forecasting* 16, 173–190.

Lo, Andrew W., 2002, The statistics of Sharpe ratios, *Financial Analysts Journal* 58, 36–52.

Lo, Andrew W., 2012, Adaptive markets and the new world order, *Financial Analysts Journal* 68, 18–29.

Meinshausen, Nicolai, Marloes H. Maathuis, and Peter Buehlmann, 2011, Asymptotic optimality of the westfall-young permutation procedure for multiple testing under dependence, *The Annals of Statistics* 39, 3369–3391.

Mills, Terence C., 1999, *The Econometric Modelling of Financial Time Series* (Cambridge University Press (CUP)).

Niederhoffer, Victor, and M. F. M. Osborne, 1966, Market making and reversal on the stock exchange, *Journal of the American Statistical Association* 61, 897–916.

Pearson, Karl, 1900, On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling, *Philosophical Magazine Series 5* 50, 157–175.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, 2011, Scikit-learn: Machine learning in python, *Journal of Machine Learning Research* 12, 2825–2830.

Pesaran, M. Hashem, and Allan Timmermann, 1992, A simple nonparametric test of predictive performance, *Journal of Business & Economic Statistics* 10, 461–465.

Phillips, Peter C. B., Yangru Wu, and Jun Yu, 2011, Explosive behavior in the 1990s Nasdaq: When did exuberance escalate asset values?, *International Economic Review* 52, 201–226.

Plackett, R. L., 1983, Karl Pearson and the chi-squared test, *International Statistical Review* 51, 59.

Romano, Joseph P, and Michael Wolf, 2005, Exact and approximate stepdown methods for multiple hypothesis testing, *Journal of the American Statistical Association* 100, 94–108.

Romano, Joseph P., and Michael Wolf, 2016, Efficient computation of adjusted p-values for resampling-based stepdown multiple testing, *Statistics & Probability Letters* 113, 38–40.

Satinover, J.B., and D. Sornette, 2012a, Cycles, determinism and persistence in agent-based games and financial time-series I, *Quantitative Finance* 12, 1051–1064.

Satinover, J.B., and D. Sornette, 2012b, Cycles, determinism and persistence in agent-based games and financial time-series II, *Quantitative Finance* 12, 1065–1078.

Senneret, Marc, Yannick Malevergne, Patrice Abry, Gerald Perrin, and Laurent Jaffres, 2016, Covariance versus precision matrix estimation for efficient asset allocation, *IEEE Journal of Selected Topics in Signal Processing* .

Sharpe, William F, 1994, The Sharpe ratio, *Portfolio Management* 21, 49–58.

Sullivan, Ryan, Allan Timmermann, and Halbert White, 1999, Data-snooping, technical trading rule performance, and the bootstrap, *The Journal of Finance* 54, 1647–1691.

Timmermann, Allan, and Clive W.J. Granger, 2004, Efficient market hypothesis and forecasting, *International Journal of Forecasting* 20, 15 – 27.

Tversky, A., and D. Kahneman, 1974, Judgment under uncertainty: Heuristics and biases, *Science* 185, 1124–1131.

Varanasi, Mahesh K., 1989, Parametric generalized Gaussian density estimation, *J. Acoust. Soc. Am.* 86, 1404.

White, Halbert, 2000, A reality check for data snooping, *Econometrica* 68, 1097–1126.

Zhang, Yi-Cheng, 1999, Towards a theory of marginally efficient markets, *Physica A: Statistical Mechanics and its Applications* 269, 30–44.

Zunino, Luciano, Massimiliano Zanin, Benjamin M. Tabak, Darío G. Pérez, and Osvaldo A. Rosso, 2009, Forbidden patterns, permutation entropy and stock market inefficiency, *Physica A: Statistical Mechanics and its Applications* 388, 2854–2864.